RX 7700 XT REVIEW PG.74

BUILD YOUR OWN **SPOTIFY** PG.58

F1 **CLOUD COMPUTING** PG.40

MAXIMUMPC

MINIMUM BS • DECEMBER 2023 • www.maximumpc.com

DO IT YOURSELF!
STEP-BY-STEP GUIDE
PG.16

BUILD A TOP-END AMD PC

- RX 7900 XTX + Ryzen 9 7900X
- Maxed out AMD for just $3,000

VRAM: HOW MUCH DO YOU NEED? PG.30

FUTURE
VOL 28, NO 12

WEB BROWSER FACE-OFF PG.50

40
DIGITALSTORM

PC GAMER
EDITOR'S CHOICE
★★★★★
MAXIMUM PC
KICK ASS!

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM

Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside

The Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm.Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

gamesradar+

One-stop destination for games, movies, TV, and more!

New Channels: Movies, TV and Cool Stuff!

www.gamesradar.com

inside

DECEMBER 2023

SCAN TO GET THE
TOM'S HARDWARE
WEEKLY NEWSLETTER

QUICKSTART

8 THE NEWS
Intel's mainstream Alchemist; Google to get AI image creation; Microsoft buys Activision Blizzard.

11 TECHTALK
Jarred Walton on the arrival of AMD FSR 3 frame generation.

12 THE LIST
The best AMD & Nvidia graphics cards.

Might there be one or two old-timers from the previous gen of GPUs?

38 SUBSCRIBE TODAY
Subscribe to *Maximum PC* and receive a digital copy of these guides, and 100+ back issues.

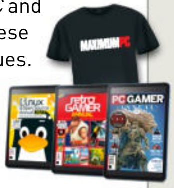

R&D

63 HOW TO
Create and export your own eBooks; Make secure backups using an ISO file.

LETTERS

14 DOCTOR

94 COMMENTS

16 TOP-END AMD PC

30 THE STATE OF VRAM
How much do you need on your GPU? We embark on a deep dive.

40 CLOUD COMPUTING
Barry Collins looks at how the cloud changed Formula One forever.

50 THE POWER BROWSERS
Bored of Chrome and Safari? Here are the best alternatives.

58 MUSIC SERVERS
Nik Rawlinson on how to run your own free music server with Alexa.

IN THE LAB

74 SAPPHIRE PULSE RX 7700 XT

82 ASUS VIVOBOOK S 15 OLED

88 PHANTEKS NV7

89 CORSAIR K70 MAX

90 TOTAL WAR: PHARAOH

© AMD, SEGA, SAPPHIRE

DISCOVER THE ALTERNATIVE TO APPLE AND WINDOWS

This is a treasure-trove of Linux and open source knowledge, from the evolution of Ubuntu and its mobile platform, to projects like automating your home with Raspberry Pi.

NEW
LINUX
2024 ANNUAL FORMAT
» HACK YOUR GRAPHICS
» GET INTO CODING
» ESCAPE WINDOWS
» FREE ALL YOUR FILES
AND MUCH MORE!
OVER 160 PAGES OF EXPERT TUTORIALS & MUST-READ FEATURES
ALL THE BEST CONTENT FROM THE NO.1 LINUX MAGAZINE
Linux Format Annual 2024
ISBN 978-1-80521-141-9
VOLUME 7
ON SALE NOW

FUTURE

Ordering is easy. Go online at:

magazinesdirect.com

Or get it from selected supermarkets & newsagents

Meet the technology experts

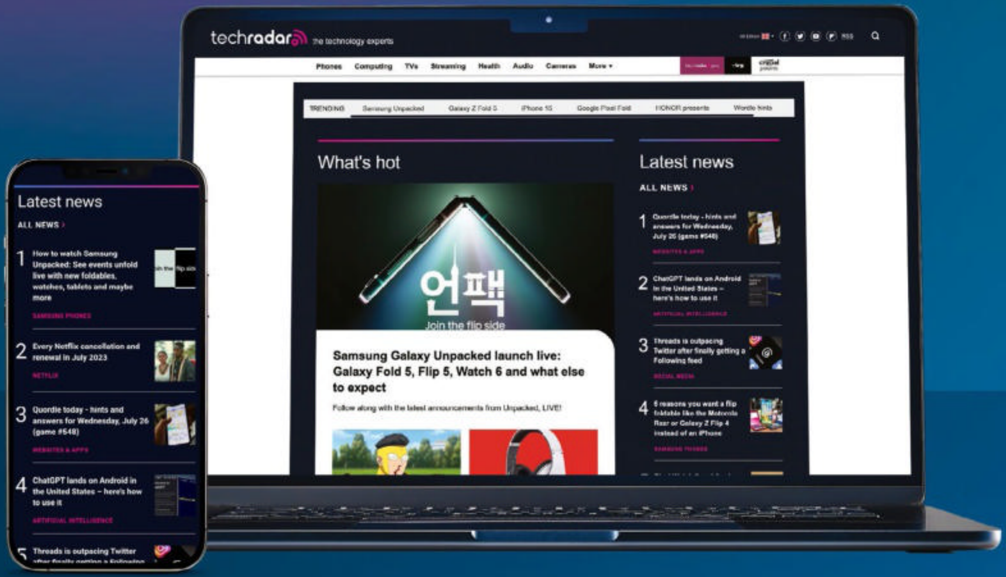
techradar the technology experts
Phones
Computing
TVs
Streaming
Health
Audio
Cameras
More
TRENDING
Samsung Unpacked
Galaxy Z Fold 5
iPhone 15
Google Pixel Fold
HONOR presents
Wordle hints
What's hot
언팩
Join the flip side
Samsung Galaxy Unpacked launch live: Galaxy Fold 5, Flip 5, Watch 6 and what else to expect
Follow along with the latest announcements from Unpacked, LIVE!
Latest news
ALL NEWS
1 Quordle today - hints and answers for Wednesday, July 26 (game #548)
2 ChatGPT lands on Android in the United States – here's how to use it
3 Threads is outpacing Twitter after finally getting a Following feed
4 6 reasons you want a flip foldable like the Motorola Razr or Galaxy Z Flip 4 instead of an iPhone
Latest news
ALL NEWS
1 How to watch Samsung Unpacked: See events unfold live with new foldables, watches, tablets and maybe more
2 Every Netflix cancellation and renewal in July 2023
3 Quordle today - hints and answers for Wednesday, July 26 (game #548)
4 ChatGPT lands on Android in the United States – here's how to use it
5 Threads is outpacing Twitter

- The world's **most** comprehensive technology website
- An **unrivalled** mix of news, opinions, reviews and features
- **All-new design**, new homepage, new features and special reports
- Backed by **over 300 years** of editorial experience

techradar
the technology experts

www.techradar.com

MAXIMUMPC

EDITORIAL
Editor-in-Chief: Guy Cocker
Staff Writer: Sam Lewis
Contributing Writers: Alex Blake, Steve Cassidy, Barry Collins, Myles Goldman, Christian Guyton, Dave James, Jeremy Laird, Rick Lake, Chris Lloyd, Darragh Murphy, Nick Peers, Nik Rawlinson, Jacob Ridley, Zak Storey, Chris Szewczyk, Jarred Walton
Production Editor: Steve Wright
Editor Emeritus: Andrew Sanchez

ART
Art Editor: Fraser McDermott
Photography: Neil Godwin, Olly Curtis, Phil Barker
Cover Photo Credits: AMD, Alamy, Arc, Sidekick, Vivaldi, Future PLC

BUSINESS
US Marketing & Strategic Partnerships: Stacy Gaines, stacy.gaines@futurenet.com
US Chief Revenue Officer: Mike Peralta, mike.peralta@futurenet.com
East Coast Account Director: Brandie Rushing, brandie.rushing@futurenet.com
East Coast Account Director: Michael Plump, michael.plump@futurenet.com
East Coast Account Director: Victoria Sanders, victoria.sanders@futurenet.com
East Coast Account Director: Melissa Planty, melissa.planty@futurenet.com
East Coast Account Director: Elizabeth Fleischman, elizabeth.fleischman@futurenet.com
West Coast Account Director: Austin Park, austin.park@futurenet.com
West Coast Account Director: Jack McAuliffe, jack.mcauliffe@futurenet.com
Director, Client Services: Tracy Lam, tracy.lam@futurenet.com

MANAGEMENT
CEO: Jon Steinberg
MD Tech: Paul Newman
Group Editor-in-Chief: Graham Barlow
Group Art Director: Warren Brown

PRODUCTION
Head of Production: Mark Constance
Senior Production Manager: Matthew Eglinton
Production Manager: Vivienne Calvert
Production Assistant: Emily Wood

Future US LLC, 130 West 42nd Street,
7th Floor, New York. NY 10036. USA. www.futureus.com

INTERNATIONAL LICENSING & SYNDICATION
Maximum PC is available for licensing and syndication.
To find out more, contact us at licensing@futurenet.com or view our available content at www.futurecontenthub.com.
Head of Print Licensing: Rachel Shaw, licensing@futurenet.com

SUBSCRIBER CUSTOMER SERVICE
Website: www.magazinesdirect.com
Tel: 844-779-2822
New Orders: help@magazinesdirect.com
Customer Service: help@mymagazine.co.uk

BACK ISSUES
Website: https://bit.ly/mpcsingleissue

Next Issue On Sale December 5th, 2023

EXTRA DIGITAL FEATURES
AUDIO FILE
PHOTO GALLERY
VIDEO FILE

© 2023 Future US, Inc. All rights reserved. No part of this magazine may be used or reproduced without the written permission of Future US, Inc. (owner). All information provided is, as far as Future (owner) is aware, based on information correct at the time of press. Readers are advised to contact manufacturers and retailers directly with regard to products/services referred to in this magazine. We welcome reader submissions, but cannot promise that they will be published or returned to you. By submitting materials to us, you agree to give Future the royalty-free, perpetual, non-exclusive right to publish and reuse your submission in any form, in any and all media, and to use your name and other information in connection with the submission.

FUTURE
Connectors.
Creators.
Experience
Makers.

Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futureplc.com

Chief Executive Officer **Jon Steinberg**
Non-Executive Chairman **Richard Huntingford**
Chief Financial and Strategy Officer **Penny Ladkin-Brand**

Tel +44 (0)1225 442 244

© GETTY IMAGES

Guy Cocker

ELECTRIC AMD

THIS ISSUE, I got back into PC building, as you can see from the RGB-laden machine adorning this month's cover. It has somehow been nearly a year since my last cover build, with our first Nvidia RTX 4090-powered machine. It seems only right that my follow-up be with AMD's highest-end card, the RX 7900 XTX. While that's also nearly a year old at time of writing, it had been an ongoing discussion between *Maximum PC* staff members that we'd somehow managed to avoid the top-end AMD card in any of our builds—it's just been a busy year for CPU and GPU releases. Thankfully, we were still impressed with what AMD has to offer, particularly with the card commonly now available with a 10 percent price cut since launch, now around $900.

In fact, one of the most interesting things for me, coming from my self-financed Intel/Nvidia system with its Core i5-13600K and RTX 4090, is how much value the AMD equivalent offers. Sure, our Ryzen 9 7900X/RX 7900 XTX build cost over $3,000, but a chunk of that went on some snazzy lighting and an ultra-fast SSD, so you could easily shave $400-500 off that. Okay, it's still a substantial chunk of change, and AMD still isn't the place to look if you want ray-traced lighting and AI-powered upscaling, but it's a system I'd still be very happy gaming on, knowing it would still be going strong for years to come. You can read all about it from page 16.

While the RX 7900 XTX has 24GB of GDDR6 memory which is more than enough to manage current gaming demands, cards in the mid-to-low end with 8GB of memory have started to feel a little restrictive of late. Our resident graphics expert, Jarred Walton, has put himself through yet more torturous testing for the benefit of *Maximum PC* readers to find out just how much of an issue 8GB of memory is in the latest titles, or if it's worth worrying about at all. Some of the results are surprising, so before you buy your next graphics card, turn to page 30 to find out more.

We also have some great tutorials, including a welcome one-off return of iFixIt's iPhone 15 Pro Max breakdown. With the spiraling costs of media subscription services, we also show you how to build your own Spotify this issue, or more accurately, stream your music collection. We also have a guide to creating your own eBooks so you can archive them and view on your Kindle reader. Don't forget to read our guide to making secure backups using ISO files that will keep your data secure for years and maybe even decades to come.

Looking ahead to next issue, we're going to be reviewing, and building with, Intel's new 14th generation CPUs, aka Meteor Lake series. We've taken delivery of the i9-14900K and the i5-14600K for a high-end and mid-range combo, which will be leading our reviews section, and we'll be building with the 14900K alongside an RTX 4090 to create the ultimate high-end machine. What's more, we have Zak Storey returning to build and write it.

Is it worth upgrading to Intel's newest generation of processors, or is it the time to look to AMD's 7000-series or Intel's 13th generation chips? And which products will feature in our coveted Gear of the Year special? Find out next issue, in our Holiday edition, on sale December 5 2023!

Enjoy the issue!

Guy Cocker

Guy is Maximum PC's *editor-in-chief. He built his first gaming PC in 1997 to play* Tomb Raider *on 3dfx, and has been obsessed with all things PC ever since.*

↘ submit your questions to: editor@maximumpc.com

THE NEWS

Intel's mainstream Alchemist

Decent 1080p gaming for $179

THERE WAS a frisson of excitement when Intel started making discrete graphics cards. At last, there would be a third player, and the possibility of a market shake-up. Maybe it would even help keep the price of high-end cards down. Intel, however, had other ideas. It started slowly, at the bottom of the market. The drivers proved under-developed, and the range started well away from the big players, whose monster cards have remained expensive. But, slowly and steadily, Intel has been building on its start. First, we got the Arc Alchemist A770, A750, and A380. Its latest offering is the Arc Alchemist A580, the first '5' series card, which has been delayed after being announced over a year ago.

The first thing to know is that it costs $179. It has been a while since we've seen a respectable graphics card launching at such a price. It's aimed squarely at 1080p gaming, and is built around a cut-down ACM-G10 GPU, the same chip as the A750 and A770. It has 24 Xe cores, and 24 ray-tracing cores. This puts it closer to the A750, which has 28 of each, than the A380, which has just eight. There's 8GB of GDDR6 memory on a 256-bit bus, which is good for a transfer rate of 512Gb/s. The 1,700MHz clock is a slight disappointment, given that the rest of the Alchemist range manage over 2GHz. It has a quoted power consumption of 185W, but it reportedly draws over 210W when working hard. It even manages to draw over 40W at idle.

Intel® Arc™ A580 Graphics Cards
ASRock
GUNNIR
SPARKLE

The Arc Alchemist A580—a cheap, mostly cheerful card with a solid 1080p performance, but still cursed with wobbly drivers.

The card has been put through thorough testing, generally against its A750 big brother, and the similarly priced AMD Radeon RX 6600. Nvidia still hasn't released a 4000-series card at anything like this price. The RX 6600 makes an interesting comparison. It's a similar 8GB card with 1080p gaming in mind. It started life at $329 in 2021. It's now about $199-209, and the RX 7900 is about $239. Running *Tomb Raider* at 1080p, the A580 is a few frames shy of 100fps, which is perfectly playable. More demanding modern AAA titles drop a little. *Balder's Gate 3* tops 70fps (but only if you switch to the Vulcan API), and *Total War: Warhammer* manages about 85. *Starfield* initially proved a disaster, broken thanks to a driver problem, but it has since largely been fixed.

Switching to 1440p gaming is slower, although it can still run *Horizon Zero Dawn* in the 70s, and *F1 2022* manages to top 90 fps. Pitted against the Radeon RX 6600, it's a mixed bag—it's 25 frames better at *Tomb Raider*, but can lag behind elsewhere. Ray tracing at 1080p using *Tomb Raider* again, the A580 can nearly hit 60fps. Move to 1440p, and it manages just over 40fps—a few frames over the RX 6600. It's a solid 1080p card, and even good for jaunts into 1440p and ray tracing.

Given the similarity in performance between the A580 and A750, Intel has dropped the price of the latter to $189. The A750 card has already come down from its launch price of $289. Given that it's just $10 more, the A750 is better if you plan to go team blue. We couldn't unreservedly recommend the A580, which isn't down to the hardware, but the shaky state of the drivers. Each new driver brings improvements though—the latest boosts some DX11 games by significant amounts (18 percent for *Counter Strike 2*, for example). Intel's cards are getting better, but still need too much fiddling. The A580's day will come though, given better drivers, and maybe a small price shave, and Intel will be onto a winner.

Intel is playing a long game—it has been working on the project for years, with three more GPU architectures in the pipeline. As it moves upmarket, it'll put pressure on prices. This was always Intel's plan: not to dominate with extraordinary performance, but to provide a better card at the price. This assault on the graphics card market is taking shape. If things continue in this vein, Intel will squeeze AMD and Nvidia where it hurts: the profitable mid-range. **–CL**

© INTEL

HIGH-SPEED STORAGE

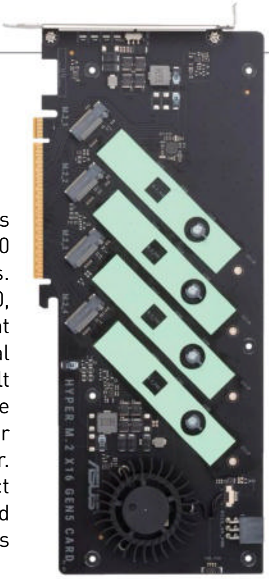

TO MAKE A PC FEEL FAST, you need fast storage, and SSDs on the processor's PCIe lanes are the way to go. PCIe 5.0 is rapidly replacing PCIe 4.0, with double the transfer rates. Asus has upgraded its Hyper M.2 x16 card to PCIe 5.0, bringing its top transfer speed to 512Gb/s (64GB/s). That figure is for four fast drives in a RAID 0 array. It's a theoretical figure, and given that M.2 drives that fast haven't been built yet, unachievable. The card has four M.2 slots, that can take 2242 to 22110 size cards. Cooling is provided by a full-cover heatsink with a fan, and power by a 6-pin PCIe connector. Some research is needed before you invest, as for full effect you'll need 16 PCIe lanes spare, so check your motherboard and processor capabilities. The Hyper M.2 x16 Gen5 costs $80, and gives a relatively easy performance upgrade. **–CL**

GOOGLE TO GET AI IMAGE CREATION

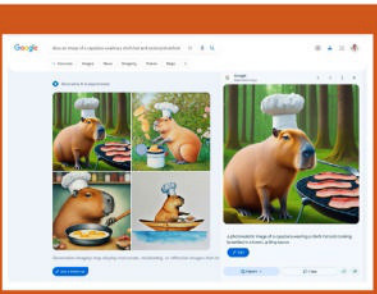

GOOGLE CALLS its AI search feature SGE (Search Generative Experience). The latest update includes the ability to generate images from text prompts, the trick that made DALL·E so famous. Your prompt will generate four images. It's also set to be part of the image search function, enabling you to create an image. There will be a filtering process—a 'strict' one, according to Google. It will also block photo-realistic faces and prompts that name notable people. The AI images will include embedded metadata that identifies their AI origins, and an invisible watermark. This won't stop the feature being misused—the 'guardrails' that AI systems have are proving ineffective. A group from Princeton University have discovered that safety systems can be side-stepped by fine-tuning the model or text strings. Human nature has a destructive streak, and AI systems are prime targets. **–CL**

Tech Triumphs and Tragedies

A monthly snapshot of what's good and bad in tech

TRIUMPHS

GOODBYE CHEMICALS
A portable system has been developed that can destroy most PFAS substances.

HOUSES ON THE MOON
NASA looks to return to the Moon, and has plans to create 3D printed houses on arrival.

TRUE BIONIC HAND
Researchers have created a prosthetic hand that is connected to nerves, enabling a test subject to achieve 80 percemt of a hand's capabilities.

TRAGEDIES

MOST NFTS WORTHLESS
A study of the NFT market by dappGambl has found that 95 percent of them are worth $0.

WIN11 NOT THE FAVORITE
Win11 is still less popular than Win10. The numbers are 24 percent against 72 percent.

WHY STADIA FAILED
Leaks have revealed why Google's foray into streaming games failed. Most tellingly, it didn't understand how much AAA game development costs.

© ASUS, GOOGLE, AMD

ZEN 5 AND 6 LEAKS

Core counts set to double, twice

FIRSTLY, this is a rumor—one that comes with a set of decidedly official-looking AMD slides apparently snapped from a monitor, and from a reliable source. This isn't gospel, but it shouldn't be far from what's coming either, and is good news. Zen 5 will be a complete architectural overhaul, although the chip structure remains the same. The branch prediction has been redesigned, and gets zero bubble conditional branches (a system that handles a branch without delaying subsequent instructions). There's more L1 cache, 50 percent bigger at 48KB. The throughput has been worked over, as has the scheduler and prefetch. All this rejigging is said to result in an IPC improvement of 10 to 15 percent.

AMD has traditionally been conservative with its IPC predictions. Zen 4 was supposed to give a bump of about eight percent, but it managed about 14. Current estimates put the actual increase at between 15 and 26 percent. We also get more cores. It's expected that there will be two CCX designs: a 4nm version with eight full Zen 5 cores, and a 3nm one with 16 Zen 5c efficiency cores. One of each would give us a 24-core range-topper. There's no news of V-Cache versions, but AMD will no doubt keep that trick for when it needs it.

A few details on Zen 6 have emerged, although more sketchy, but do indicate another doubling of the core count. There's also an FP16 unit for AI work. The estimate here is for an IPC hike of at least ten percent. Importantly, it will mark a move away from the current Zen layout with an I/O chiplet, and one or two CCXs that we've had since Zen 2. Instead, we'll have a new generation of the infinity fabric, and new packaging techniques. It'll use both a 3nm and 2nm process for its CCXs. The first Zen 5 chips are expected next year. It'll go against Intel's Arrow Lake, which looks good, but it'll be along later, and looks to be relatively expensive and power-hungry. Even given Intel's recent improvements, AMD is still looking extremely strong. **–CL**

MICROSOFT BUYS ACTIVISION BLIZZARD

The $68.7 billion deal is finally made

THE REGULATORS in the UK have acquiesced, and the deal to buy Activision Blizzard is on. The chief concern for regulators was always that Microsoft would have too tight a grip on cloud gaming. More specifically, it would be able to make *Call of Duty* (the fourth largest game franchise in history) exclusive to the Xbox and Windows.

The competition regulators did squeeze concessions out of Microsoft. First, it had to agree to keep *Call of Duty* available on PlayStation for ten years. Also, the UK's regulators made the company sell exclusive worldwide rights and non-exlusive rights for the European Economic Area to Ubisoft for the next 15 years. According to Microsoft's vice chairman, Brad Smith, this "addresses the concerns about the impact of the proposed acquisition on cloud game streaming". At least all parties are now being polite.

Microsoft now has a veritable treasure chest of top-tier games, including *World of Warcraft*, *Diablo*, *Candy Crush*, *Overwatch*, *Hearthstone*, *Starcraft*, and more to add to its Xbox stable. Titles will be popping-up on Game Pass quickly now, although reportedly *Diablo IV* and *Modern Warfare 3* won't make it this year.

The concessions aren't as onerous as they seem: it still gets the lion's share of any in-game purchases, which is a lucrative part of gaming. There are also time limits, so eventually it'll own the rights entirely. The jury is out on whether this is a good deal for gamers. It is the biggest deal in gaming history, and Microsoft has a massive investment to earn back. It's expected to push players to Game Pass, for starters. The FTC hasn't completely given up on opposing the deal either—it has a date with the appeals court in December, but it looks highly unlikely to be able to do much now. Welcome to your new home, Activision. **–CL**

GeForce RTX 4080 Ti coming

EARLY IN THE NEW YEAR, we're expecting Nvidia to release the GeForce RTX 4080 Ti. The regular RTX 4080, while awesome, was deemed too expensive, with a launch price of $1,199. It was expected to follow the established pricing pattern for GeForce cards, and significantly undercut the top-dog RTX 4090 (which launched at $1,599). The previous generation's RTX 3080 started at $699, so even allowing for inflation, it was an unexpected move. It also generated an uncomfortable gap in performance, with the RTX 4080 having 9,728 shaders, while its big brother 4090 had a 16,384 count for not much more money. This made the the 4080 look particularly poor value, and left a huge performance gap in the 4000-series range.

Adding an RTX 4080 Ti gives room to drop the RTX 4080's price by $100-$200, and even out the pricing structure at the top end. It would also let the 4080 compete with AMD's Radeon RX 7900 XTX, which is as good at rasterized 1440p gaming, and $150 or more cheaper (see this month's build on page 16). The RTX 4080 Ti is expected to host a AD102 GPU, with 14,080 shaders, run to around 450W, and sell for $1,100-1,200. This would put it closer to a 4090 than the 4080 in performance, but much depends on memory and bus speed. The RTX 4080 was never popular, as people knew they were being sold a dud. You either spent the extra, and went big with the monster RTX 4090, or jumped ship to AMD. Nvidia needs to get this right. **–CL**

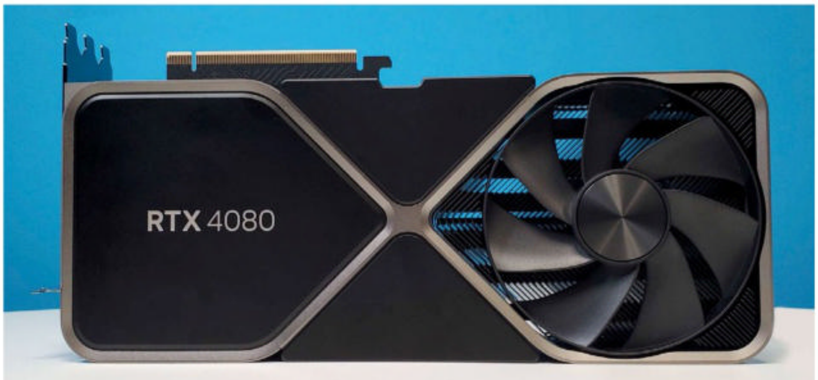

Nvidia's unloved RTX 4080, where the price/performance balance was terrible. A new Ti version should fix this.

16-pin problem

Nvidia introduced the 12VHPWR socket to get enough juice to the GeForce RTX 4090. It averages a draw of around 380W while gaming, and can peak at over 600W. There were problems from the start, with overheating cables and the occasional fried card. After blaming user error, Nvidia redesigned the connector, dubbed the 12V-2x6. Things quietened down, but not completely. A new failure has emerged in a card that reportedly had been working perfectly for a year. There are still a number of 4090's in circulation with the original connector that could fail, but at least they'll be under warranty. CableMod has just released a revised version of its adapter that claims to fix stress and cooling issues, all for $39.90. **–CL**

Meteor Lake desktops

Intel's Meteor Lake won't be a desktop part, as there won't be socketed versions; it'll be ball grid array packaging only. Due to technical issues, it is to be aimed at mobile platforms. However, it'll make its way into desktop designs. A company statement claims it "will power innovative mobile and desktop designs, including desktop form factors such as All-in-One". One reason seems clear from the early benchmarks of engineering samples. An Ultra 7 165H (16 cores, 22 threads, and a 5GHz top clock) has been run through Geekbench, and was as fast as a Core i7-13700H. Nice, but not serious desktop material. We'll have to wait until the end of next year, and the first Arrow Lake chips. **–CL**

© ACTIVISION PR

Jarred Walton

TECH TALK

AMD FSR 3 Frame Generation Arrives

AMD'S FIDELITYFX SUPER RESOLUTION 3 (FSR3) technology was first announced in November 2022—a clear attempt to show that AMD could also provide frame generation technology. One year later, we have the first two samples of how FSR3 works and how it looks. It's not a great first outing. There are many issues, though AMD seems aware of them, and apparently has fixes in the works.

This raises the question of why it would choose two less than ideal examples as a launch showcase for the tech: *Forspoken* and *Immortals of Aveum*.

Forspoken gets a strike against it thanks to its 120 fps framerate cap (240 fps with Frame Generation enabled). We've had 144 Hz and 240 Hz monitors for many years, with 480 Hz and even 540 Hz displays now available. There's no way to take advantage of such hardware with *Forspoken*.

Immortals of Aveum doesn't have an artificial framerate cap, but there are other issues with FSR3 frame generation. In testing, it often felt jittery and choppy, but then we get to the strange requirements. FSR3 frame generation has known problems with variable refresh rate technologies, and AMD recommends turning on vsync.

These apparently provide more consistent frame pacing, but frame generation, as implemented in FSR3, breaks some display features. We would almost always recommend using a VRR display where possible, as it helps to reduce latency. This is particularly true with games that are relatively demanding, as many GPUs won't be able to hit higher framerates. The RTX 4090 only averaged 92 fps, with FSR3 frame generation enabled. On a 120 Hz or higher G-Sync compatible display, you'll get a worse experience with frame generation on.

But even with the appropriate settings, in our initial testing, minimum framerates with frame generation were worse than the base performance. An RTX 4060 Ti ran *Forspoken* at 100 fps on average, with one percent lows of 79 fps. Turning on frame generation bumped the average to 144 fps, but one percent lows dropped to 70 fps. That meant more perceptible stuttering, and as resolution increases, this gets more problematic.

Still images look fine in *Forspoken*, but the results can be a bit less consistent when in motion.

© SQUARE ENIX

As with Nvidia's DLSS 3 frame generation, FSR3 interpolates between two rendered frames, delaying when you'll see the second fully rendered frame. That means higher latency. Conveniently, both games support DLSS upscaling. FSR3 frame generation is also locked to use with FSR3—it's not possible to use DLSS upscaling with FSR3 frame generation.

Finally, there's image quality. DLSS 3 frame generation isn't perfect, but the FSR3 generated frames were more prone to issues. Certain areas in *Immortals of Aveum* showed horrible shimmering and other artifacts. Finely detailed areas also tended to have issues, more noticeable at either lower fps, or if we captured the output and looked at individual frames. Any moderately fast movement tended to look less than ideal, and FSR3 turns off frame generation completely if the frame to frame difference is too high.

At least FSR3 is now available, and works with most GPUs. Several AMD and Nvidia cards worked okay, though while Arc A750 also worked, the Arc A770 crashed to desktop every time we tried to enable frame generation. Regardless, FSR3 frame generation is far from a great experience, and like DLSS, it needs high base framerates before it's viable. That means the GPUs that could benefit most from it—slower cards from a generation or two back—end up being a poor experience. Hopefully future updates and games will improve the feature and at least allow for VRR with FSR3 frame generation.

Jarred Walton has been a PC and gaming enthusiast for over 30 years.

THE LIST

THE BEST AMD & NVIDIA GRAPHICS CARDS

LIKE A LOT OF PC COMPONENTS, graphics cards are cyclical. We're now in the calm between the storms of new GPU generation launches. Both Nvidia's RTX 40 Series and AMD's RX 7000 GPUs have been fully fleshed out. Likewise, it'll be a while before any properly new graphics cards are launched, so it's the perfect time to take stock and consider where the value propositions lie. In other words, what are the best picks of this generation? And might there be one or two old timers from the previous gen of GPUs that still make sense?

5 AMD RADEON RX 7600

There's nothing we'd enthusiastically recommend at the budget end, but the best of a bad bunch is the AMD Radeon RX 7600, especially now it's $239. Like all AMD GPUs, it's light on features, and ray-tracing performance isn't great. But for affordable 1080p gaming, it's solid. Moreover, it comes close to the RTX 4060 for regular raster performance at a lower price. **www.amd.com, $239**

4 AMD RADEON RX 6700 XT

Another old-timer from AMD. But for a whisker over $300, the 6700 XT is now a great deal. For sure, it's light on features compared to Nvidia's latest GPUs. Ray tracing is a conspicuous weakness, and AMD's new frame generation tech has yet to be proven. But for pure raster performance, you're getting a lot of bang for a very reasonable amount of buck. It'll murder the RTX 4060 in most games, that's for sure. **www.amd.com, $309**

3 NVIDIA GEFORCE RTX 4070

At launch for $599, the RTX was hard to stomach, but this is a good all-round GPU for the money. It performs close to the previous-gen RTX 3080 in many games, and adds Nvidia's strong feature set. The main drawback is the mere 12GB of VRAM and a narrow memory bus. **www.nvidia.com, $529**

2 AMD RADEON RX 6800

AMD's RX 7000 Series didn't bring major advances with features like ray-tracing, so the fact that this card uses AMD's flagship GPu from the previous gen means it can outperform AMD's new RX 7700 XT. The fact that it has a nice, big 256-bit memory bus and 16GB of VRAM means that it's likely to stand the test of time better than the 192-bit 12GB 7700 XT, too. **www.amd.com, $419**

1 NVIDIA GEFORCE RTX 4090

Okay, it's megabucks. But it's the only member of Nvidia's RTX 40 Series that delivered a big step in performance for little to no price increase over its predecessor. In that sense, it's the best-value RTX 40 GPU of all. It also happens to be the fastest graphics card you can buy. It's as close as it gets to future-proofing when it comes to GPUs. **www.nvidia.com, $1,599**

RTX 4090
GEFORCE RTX

© AMD,NVIDIA

Jeremy Laird

TRADE CHAT

Is it game over for AMD graphics?

SPARE A THOUGHT for AMD's GPU engineers. Every time they think the gap to Nvidia is closing, team green comes up with some new wizardry, and it's back to playing catch up. Like 'em or loathe 'em, technologies like ray tracing and DLSS have gained major traction in the PC gaming market. AMD has struggled to respond. So how much longer can this continue?

Constantly catching up. That's what it must feel like for AMD when it comes to PC graphics. Wind back the clock to 2004, and Nvidia rolled out SLI (Scalable Link Interface). It was a tool for connecting multiple graphics cards. I never liked it—AMD (or ATI, before it was acquired by AMD) felt obliged to copy SLI with Crossfire.

Then came Nvidia's G-Sync in 2013. As a technology that had more merits than SLI, even if Nvidia's proprietary implementation undermined its appeal, AMD copied it with FreeSync. Next came the GeForce RTX 20 Series in 2018, and the introduction of hardware ray tracing, something AMD wouldn't match for two years.

Another major milestone was Nvidia's DLSS upscaling technology in 2019, albeit it arguably didn't truly arrive as a worthwhile technology until the massively improved DLSS 2.0 update in 2020. AMD's response in FSR appeared in 2021, but if anything, AMD has only fallen further behind since.

AMD's second-generation ray-tracing technology in the Radeon RX 7000, launched last year, failed to close the gap to Nvidia. FSR 1.0 was likewise a pretty slapdash effort. Granted, FSR 2.0 was better. But along came DLSS 3.0 in 2022 with not only much-improved upscaling quality, but also Frame Generation.

GEFORCE RTX
RTX 2080 Ti

© NVIDIA

> Nvidia's controversial RTX 20 series has proven a turning point for the graphics market

AMD is only now releasing its own take on Frame Generation, which is clunky and lacking in polish. But even before AMD managed to push that feature stumbling reluctantly out the door, Nvidia was unveiling DLSS 3.5 with Ray Reconstruction, further bolstering Nvidia's dominant ray-tracing performance advantage.

Meanwhile, AMD's market share in discrete PC graphics cards has tanked. For the last year or so, AMD has bounced around between 8 percent and 15 percent, with Nvidia near 90 percent and Intel taking a few scraps. I can't see how that is sustainable. AMD will no doubt retain its GPU engineering capability for the game console and machine learning markets, but they must be considering whether add-in PC graphics cards are worth the effort.

Also, it's remarkable to observe that so many last-gen AMD GPUs, such as the RX 6700 XT and both RX 6800 and RX 6900 variants, are available at ever lower prices. AMD even recently rolled out a "new" RX 6750 GRE GPU for some markets, nearly a year after the first RX 7000 GPUs were made available.

All the while, the Radeon RX 6000 series barely registers in the latest Steam surveys of gaming hardware, despite going on sale in late 2020. It's hardly a surprise, then, that AMD was slow to fill out the current RX 7000 series with the 7700 XT and 7800 XT, which have only just gone on sale, what with piles of older GPUs still to shift.

More recently, rumors that AMD has decided to ditch the high-end variants of its next-gen RX 8000 family have been circulating. Frankly, that seems plausible. Admittedly, there's also talk that AMD will take a tilt at the performance crown with the GPU generation that follows—call it the RX 9000 series if you like.

Maybe that's true. If it is, I predict that it will be AMD's last-ditch attempt to have an impact on the discrete PC graphics market. It's clear that AMD's chiplet tech in the RX 7000 series hasn't worked as planned. Assuming that AMD tries to get back on top in raw performance terms after taking a timeout for the next generation, if the 9000 series fails, it'll be game over for high-performance AMD graphics on the PC.

Six raw 4K panels for breakfast, laced with extract of x86... Jeremy Laird eats and breathes PC technology.

DOCTOR

THIS MONTH THE DOCTOR TACKLES…

> Error-correcting RAM
> Stream eBooks
> Half-broken desktop

ECC server memory

I'm looking at building a home server using one of the boards you recommended in your September 2022 issue, but none of them support ECC memory. I understand that this is important for detecting potential memory corruption before it can lead to serious data loss. Is there a Pentium J-series board (or similar low-powered board) that supports ECC memory?

—Graham Naylor

THE DOCTOR RESPONDS: ECC stands for Error Correction Code, which refers to memory that constantly checks for errors in data memory that can indicate failing RAM. In turn, these errors can start to silently corrupt files—including personal data, as well as those from the underlying operating system. ECC memory is usually found in high-end servers in professional and business environments, where corruption or service outages could be catastrophic. Sadly, you won't find ECC memory support in any Pentium J-series chips.

You might want to examine products based on the Intel Atom Processor C series, which are designed for microservers, but there's a premium to pay—not just for the Atom C, but for ECC RAM itself. While there are some consumer-grade motherboards and CPUs that support ECC RAM, it's not a must-have component.

If you follow our best-practice advice to ensure your data is regularly backed up using differential or incremental backup tools to retain multiple versions of a file, then in the unlikely event that the RAM does start corrupting your data, there's a version you can roll back to. Duplicati (www.duplicati.com) is a good option, which we featured in the May 2023 issue.

It's also worth investing time post-server setup into creating a RAM-testing regime. Two apps can help—Memtest86+ (www.memtest.org) and Memtester (https://pyropus.ca./software/memtester). Install them both with:

```
sudo apt update && sudo
apt install memtest86+
memtester
```

Memtester can be used from the Terminal while your server is up and running, but is only capable of testing memory that's marked as free. Use the 'free -h' command to determine how much RAM is currently free, then run Memtester as follows (the example will test 4GB RAM over three iterations):

```
memtester 4G 3
```

Wait for the test to complete, and then you'll see a list of test results. Hopefully everything is 'ok'.

For a more thorough test, run Memtest86+, which can only be accessed outside of Ubuntu Server (it installs itself to your GRUB boot menu). Once a month or so, connect your server to a keyboard and monitor, then reboot and launch Memtest86+. Let it run all tests at least once—preferably up to five times. If no errors are detected, press Esc to reboot back into your server proper.

If errors are found, you'll need to test each stick of RAM by shutting down your server and removing each one in turn, then re-running the test until you identify the failing stick. It's technically possible to keep faulty RAM running—some errors are marginal and may never result in corruption, while you can also instruct Linux to avoid using bad memory areas. However, it's not simple task—see https://askubuntu.com/questions/908925 for details—and your best course of action is to simply replace the bad RAM.

Self-hosted eBooks

I'm looking for something to rival Audiobookshelf, but for eBooks, so I can stream my eBooks collection to my phone or tablet (I don't have a Kindle). I have a library of hundreds, so it's not practical to store them all on my mobile.

I've tried the web server component of Calibre, but while functional, it's not exactly attractive, and the app's full text search capabilities haven't made it to the web server yet. Is there a good alternative?

—Henry Perkins

THE DOCTOR RESPONDS: The Doc has recently discovered Kavita (www.kavitareader.com), which ticks all your boxes—including the full-text search—and boasts a slick user interface that's not dissimilar from self-hosted audiobook streamer Audiobookshelf. You can install it natively across

↘ submit your questions to: doctor@maximumpc.com

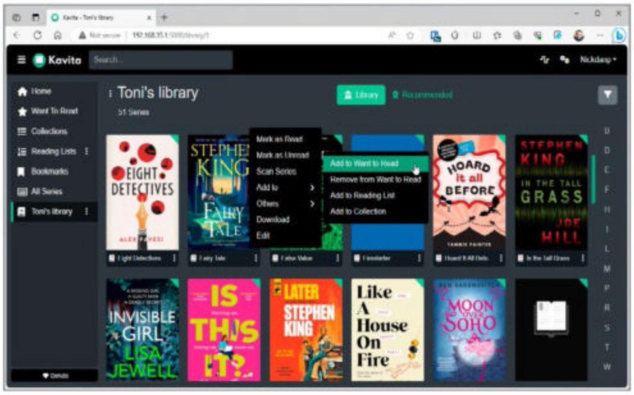

Set up your own eBook-streaming solution with Kavita.

© KAVITA

all three major platforms (Windows, macOS, and Linux), but there's also a Docker instance, which works perfectly with a home server setup, and is easy to link up to your reverse proxy for remote access.

It utilizes the same folder structure as Calibre (author/title), so assuming your books are organized, point Kavita to the folder housing your Calibre library, and it should pick everything up. A subscription unlocks some extra goodies, but the majority of functionality is free—including access to your own server's content.

The only downside is the lack of an official mobile app, but the web interface is optimized for mobile browsers, and it has built-in support for OPDS, so you can connect to it through third-party apps like those listed on Kavita's wiki (https://wiki.kavitareader.com/en/faq/external-readers).

Excel annoyance

Whenever I go to save a specific Excel spreadsheet, I'm told it may have been changed by another user, and am asked if I want to create a copy or overwrite any changes. I'm not aware of opening this anywhere else, so how do I stop the prompt from constantly appearing?

—Lynn Monson

THE DOCTOR REPLIES: We examined Lynn's file, and saw that it contained links to other workbooks that threw up a security warning and exhibited non-consistent behaviors—sometimes the file updated them with no complaint, at others it prompted the user to do so (then claimed it couldn't update all the links). Occasionally, it gave a warning about disabling automatic link updates.

We started by clicking Edit Links… when told certain links couldn't be updated. This opened a dialog revealing links to missing files. We spoke to Lynn, who confirmed the files had been moved elsewhere—by clicking 'Change Source…' next to each missing link, she was able to restore the link manually.

This resolved that error message, but there was still the security warning and claim the file had been edited by another user. We attempted to fix the former by clicking the 'Automatic update of links has been disabled' to be taken to the file's Info tab in Excel, where a Security Warning contained an 'Enable Content' button that added the document to Excel's Trusted list. But the setting didn't stick.

It turned out she was a Nextcloud user, and the file was being synced to a self-hosted Nextcloud instance (version 27.1.0). We asked her to move the file to a folder not monitored by Nextcloud, and the prompt ceased. We looked online and found other users had suffered a similar problem. Lynn had recently updated her Nextcloud client to v3.10, while online it was claimed on both Github (https://github.com/nextcloud/desktop/issues/3456) and the Nextcloud forum (https://help.nextcloud.com/t/157229/21) that the problem had been resolved as of version 3.9.2.

Clearly this isn't the case, but in both cases, Nextcloud has been identified as the culprit through its inability to record the microseconds component recorded in a file's modified date, which leads Excel to conclude that the file has been changed by someone else. While Lynn will have to put up with this for now, the fact it's been identified as a problem has us hoping Nextcloud resolves it.

No webcam audio

After a recent update, my C920 Logitech Pro Webcam's audio no longer works when I plug it in. Instead I'm told "no audio device is installed". Can you suggest a fix?

—Salvo Ramirez

THE DOCTOR RESPONDS: The Doc managed to track down a workaround. The cause is apparently a conflict between the Logitech Webcam and a particular patch that targets Realtek's AudioProcessingObject, which results in the Windows Audio service crashing.

To circumvent the problem, press Win + R and type devmgmt.msc to open Device Manager. Expand 'Sound, video and game controllers' to reveal the HD Pro Webcam C920 audio driver. Right-click this and choose Disable, confirming any prompts. Once done, audio should be restored in around 30 seconds; if not, open Task Manager, select Services, then Audosrv (which should have Windows Audio as its description), and click the Restart button.

The C920 audio driver should remain disabled until you restart it.

Desktop split personality

It appears that half of my desktop no longer works. It's visible on-screen, but I cannot interact with anything on the left-hand side of the screen, nor drag items over from the right to the left. I'm fully up to date and tried running sfc/scannow with no success. Please help!

—Mary D Mathis

THE DOCTOR RESPONDS: Mary is right when she says something invisible is getting in the way. A search revealed the culprit—and temporary workaround: Windows 11 widgets. We instructed Mary to open Task Manager by right-clicking the Taskbar, then told her to type widget into the search box at the top to reveal a Windows Widgets entry. Clicking > next to this revealed two running processes—we told her to close widgets.exe, and the problem vanished.

This doesn't survive a reboot, but if you're not using the widgets feature, you can disable the Widgets button from the Taskbar to prevent the invisible window from appearing—navigate to Settings > Personalization > Taskbar to flick it to off. If this doesn't work, disable widgets via the Registry. Press Win + R, type regedit, hit Enter, then navigate to HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft.

Right-click the Microsoft key and choose New > Key. Name it Dsh. With Dsh selected, right-click inside the right-hand pane and choose New > DWORD (32-bit) Value. Name this AllowNewsAndInterests, and leave its value set at 0 (zero).

Open Task Manager, locate the Windows Explorer process, right-click it, and choose Restart. The problem should now be fixed. ⏻

INGREDIENTS

PART		PRICE
CPU	AMD Ryzen 9 7900X	$449
GPU	AMD Radeon RX 7900 XTX	$899
Mobo	ASRock X670E Pro RS	$272
RAM	NADATA XPG LANCER RGB 32 GB (2 x 16 GB) DDR5-7200 CL34	$140
Storage	Crucial T700 2TB Gen5	$297
Case	Phanteks NV7	$200
PSU	NZXT C1000 80+ Gold	$155
CPU cooler	NZXT Kraken Elite X73 RGB AIO	$300
Fans	Phanteks D30 (3x3-pack @ $90 each, 2x1-pack @ $30 each)	$330
OS	Windows 11 Home 64-bit OEM	$32
TOTAL		**$3,074**

PRICES CORRECT AT THE TIME OF PUBLISHING

TOP-END AMD PC

With the RX 7900 XTX going for around $900, is now the time for a Team Red 4K gaming machine? *Guy Cocker* finds out.

WE KNOW, WE KNOW—for the last few issues now, we've been infatuated with all things AMD. Part of it is down to the lack of new Intel hardware coming through (which won't be a problem next issue, as we already have Raptor Lake in-hand), and part of it is just the pure excitement of what AMD's been doing of late in bringing the fight to Nvidia in the graphics card space. But somehow, despite all of this AMD love, we've still managed to avoid doing a build with the company's top-end offering, the RX 7900 XTX. Well, we're very pleased to address that oversight here.

While the RX 7900 XTX is nearly one year old at time of writing, there's no heir apparent that we're aware of in the works from AMD, nor a competitor product from Nvidia at the $1,000 mark. Officially anyway. What has happened is that Nvidia seems to have found a limit to the number of people willing to spend four figures on a GPU, so its RTX 4080 has started to drop from its $1,200 list price and can now be found at $999, at least in the form of Zotac partner boards. As a result, the RX 7900 XTX has also had its own discounts, going down $100 from $999 to $899, which while no bargain, is at least a 10 percent saving and keeps it competitive against Nvidia's rival.

But in late 2023, with Nvidia's ray tracing and upscaling technologies now generations ahead of AMD's, is it still worth investing in an RDNA 3-based card? Or is AMD's rasterization performance and value proposition just too good to miss? That's what we intend to answer with this RX 7900 XTX build. As a nice little bonus, this is also the first PCIe 5.0 build we've ever done thanks to us having access to one of Crucial's brand new SSD drives, theoretically capable of 12,400MB/s read speeds. Anyway, enough waffle—get your Phillips head screwdriver at the ready and join us as we build a top-end AMD gaming machine. –GUY COCKER

BUILD IT!
STEP-BY-STEP GUIDE
PG.22

CPU

AMD RYZEN 9 7900X

WWW.AMD.COM

Truth be told, we'd loved to have had access to the $700 Ryzen 9 7950X3D CPU for this build just to be able to see what an absolutely maxed-out AMD system could do. Especially as we already used the X3D CPU in our June 2023 cover build, where we paired the best AMD CPU with the best Nvidia GPU, the RTX 4090, for what was (and still is) the fastest gaming PC money can buy. But in practice, we would probably urge most gamers to spend on the GPU and save some money going for a more modest CPU that will offer most of the performance for a lot less cash. The Ryzen 9 7900X is still a monster CPU, with 12 cores and 24 threads, plus on-board Radeon graphics—not that they're needed here.

• Price: $449

BUILD IT!
STEP-BY-STEP GUIDE
PG.22

GPU

AMD RADEON RX 7900 XTX

WWW.AMD.COM

The rock star of this build is AMD's top-end GPU. We're using the reference board, as that's what we had to hand, but don't be afraid to shop around for partner boards, especially if you find one on offer for under $900. There's a lot to like about this GPU, which we reviewed in the February 2023 issue, and awarded a score of 8. First of all, there's a mammoth 24GB of GDDR6 memory. Then there's the fact that it needs only two power cables and no fiddly adaptors in order to run. It doesn't require a support bracket to bear its weight either, unlike Nvidia's chunky top-end cards. The only real downside is that it's a power hog at 355W, meaning AMD recommends a 800W PSU minimum to run it.

• Price: $899

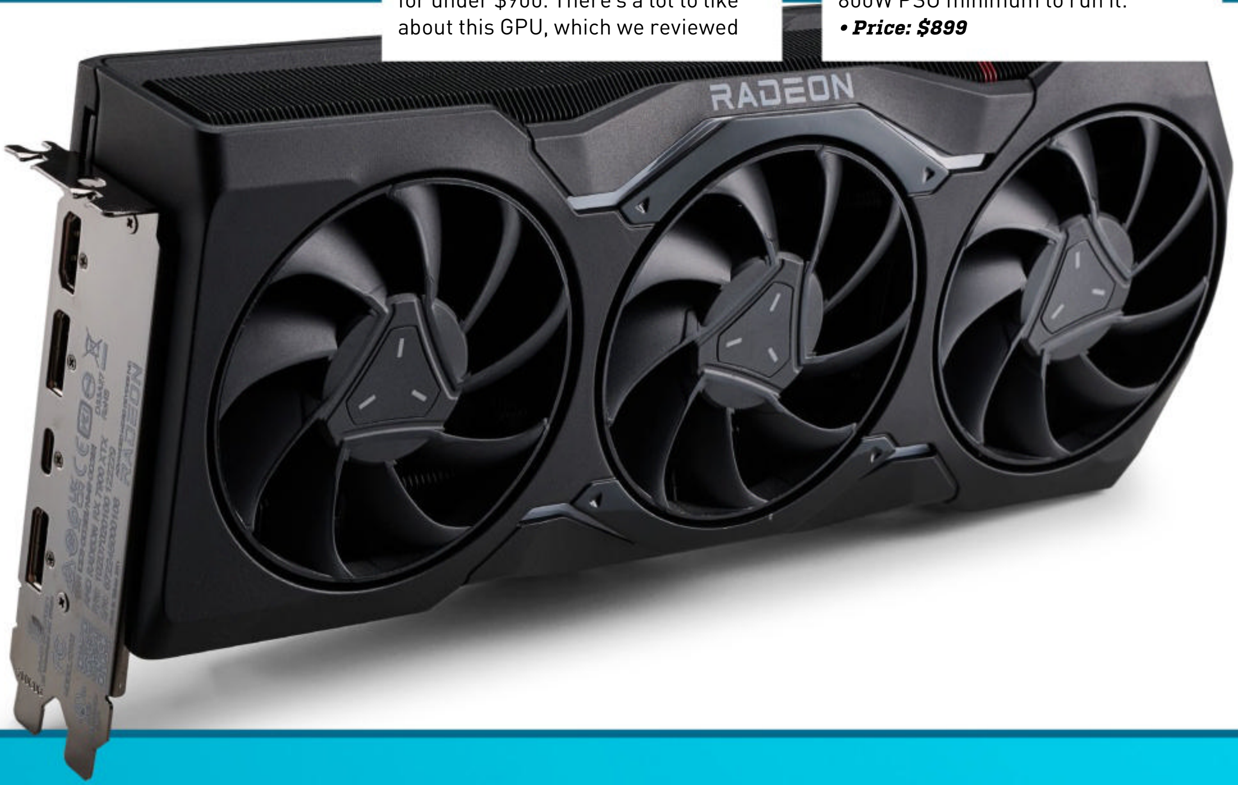

STORAGE

CRUCIAL T700 2TB GEN5

WWW.CRUCIAL.COM

While we said the RX 7900 XTX was the star of the show, this PCIe Gen 5 SSD from Crucial comes a close second. In the real world, we'd probably advocate going for a 2TB PCIe 4.0 drive for half the price, or a 4TB drive for around the same price as this. But the cutting-edge storage fiends among you, we couldn't resist pairing the compatible ASRock X670E Pro RS motherboard with the latest generation of storage. Hopefully, we'll get our fastest-ever read and write benchmarks for this drive, with quoted performance 1.8 times that of the previous gen. This drive also comes with a five-year warranty, plus a heatsink, but as our ASRock board comes with a heatsink, we've used a bare drive.

• Price: $297

crucial
by Micron
T700
PCIe® Gen5 NVMe™ M.2 SSD

MOBO

ASROCK X670E PRO RS

WWW.ASROCK.COM

We don't tend to use a lot of ASRock motherboards in our builds, so we were very happy to receive one of their boards for this machine. Especially as our sister title PC Gamer recommends it as the best affordable X670E motherboard around, at a very wallet-friendly $272. While its black design does contrast with the white finish of our other components, you don't really need to care about looks when you're getting such astonishing tech, including PCIe 5.0 GPU and SSD support, the ability to run up to a Ryzen 9 7950X without issue, and essentials like Wi-Fi 6E and ample storage support. ASRock even includes a graphics card support device if you decide to go for a high-end Nvidia GPU in future.

• Price: $272

FANS

PHANTEKS D30

WWW.PHANTEKS.COM

We admit it—we went overboard when it came to fans. But when we requested the case from Phanteks, they were kind enough to bundle in 11 of their new D30 fans in various configurations, so we wanted to go all out and make a great-looking system. When buying these fans, you can get them individually or in 3-packs, and in regular or reverse airflow configurations. Their main selling point is that they come with linking brackets that allow you to daisy-chain fans together without any clutter, as well as covers to hide the screws holding them in the case. It's then quick and easy to hook up the closest fan to the rear-mounted lighting controller for a uniform look. **• *Price: $330***

PSU

NZXT C1000 80+ GOLD

WWW.NZXT.COM

Another month, another NZXT power supply. In our defense, we tried to use another manufacturer's 1000W PSU, but when we connected it, nothing happened apart from a disconcerting whining noise. Fearing blowing up the system or damaging the components, we swapped the PSU and cables for an NZXT C850 850W model we found in the *Maximum PC* supply cupboard, but it also failed. Concerned, we grabbed a C1000 and its cables from a previous build (which we knew worked). Phew! **• *Price: $155***

CPU COOLER

NZXT KRAKEN ELITE X73 RGB AIO

WWW.NZXT.COM

Again, we've used this cooler from NZXT a few times now, but it keeps coming back for a reason—it's a really great AIO. Again, we initally tried to use a cooler from a different manufacturer, but the only one we could find was a few years old, dusty, and not in great shape. When you're building a 4K gaming rig, it pays to use good-quality, dependable components. While we're not planning on overclocking this build, going for a well-made cooler will allow us extra headroom to push the CPU harder if we decide to in the future. We also love the style of the infinity mirror display on this particular cooler, as it appears to go on forever. This time, we swapped out NZXT's fans in favor of ones provided by Phanteks that were designed for this very case (see below). **• *Price: $200***

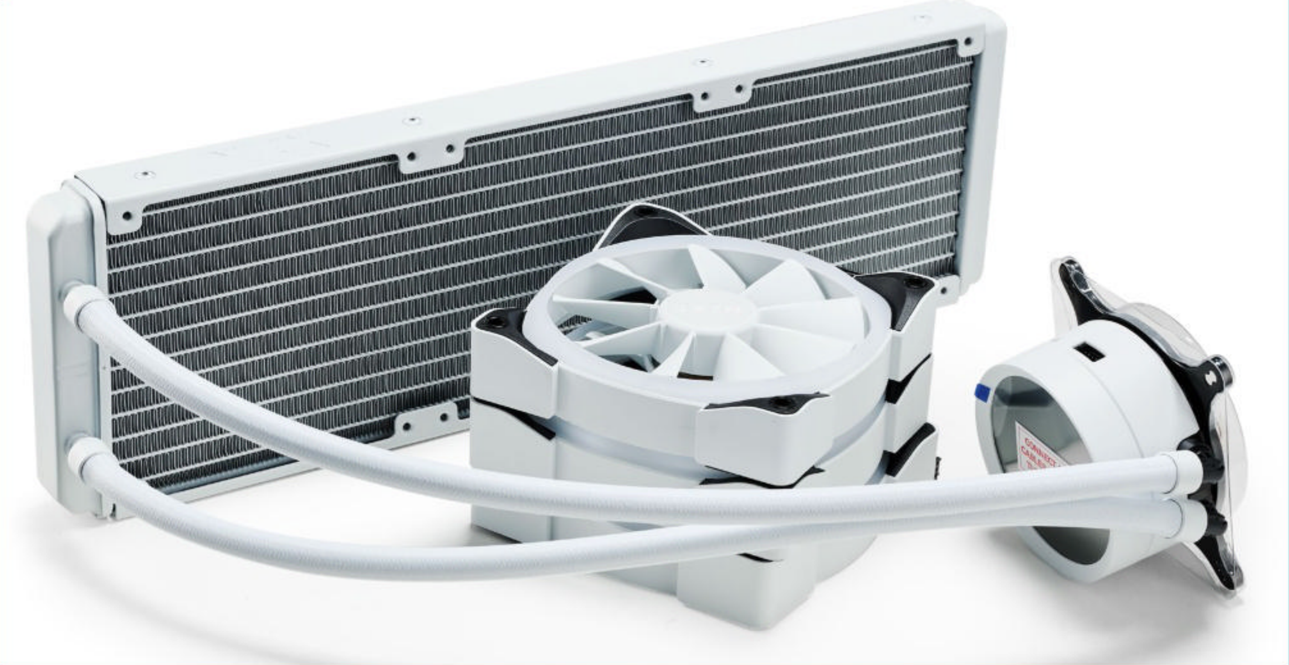

BUILD IT!
STEP-BY-STEP GUIDE
PG.22

DDR5
XPG
DDR5
XPG

RAM

ADATA XPG LANCER RGB 32 GB (2 X 16 GB) DDR5-7200 CL34

WWW.ADATA.COM

32GB of incredibly fast DDR5 RAM with RGB support for $140? Yes please! ADATA has very much understood the assignment when it comes to its XPG LANCER RAM, which our sister title Tom's Hardware has found pairs well with AMD processors. It's not the best-quality RAM on the market, using Samsung B-die over SK hynix's more overclocker-friendly ICs, but it can still easily and stably be pushed—not that you'd really want to, given the out-of-the-box speeds. It also offers great value for money, and as the photography will attest, the RGB lighting is bright and colorful. 32GB should also be more than enough for both gaming and creative work.

- ***Price: $140***

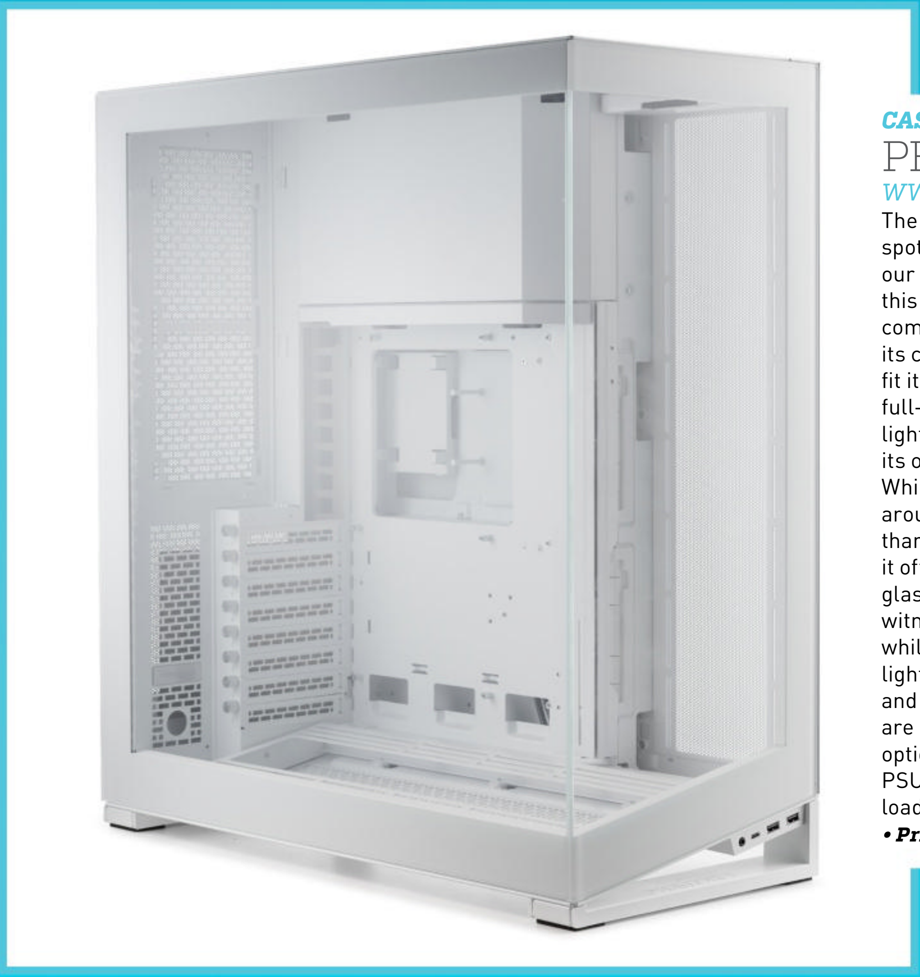

CASE

PHANTEKS NV7

WWW.PHANTEKS.COM

The eagle-eyed among you will have spotted that this case also adorned our June 2023 cover build. But this time, the Netherlands-based company also sent in a regiment of its complementary D30 RGB fans to fit it out with. This massive case is full-tower size, and it certainly isn't lightweight, coming in at 37.25 lbs on its own without any components in it. While you won't want to be moving it around too much after it's completed, thankfully you'll just want to show it off as is, as it looks awesome. Its glass corner design allows you to witness your creation in all its glory, while the white finish allows the RGB lighting to bounce off every surface and illuminate your space. There are loads of different fan-mounting options, a nice vertical mount for the PSU above the motherboard, and loads of panels to hide cabling.

- ***Price: $200***

ALL-OUT AMD

BUILD AN ALL-AMD GAMING POWERHOUSE

LENGTH OF TIME: **3-4 HOURS**

LEVEL OF DIFFICULTY: **MEDIUM**

1 Gut the PC

The Phanteks NV7 may be the biggest and roomiest case that I've ever built in for *Maximum PC*, but it's still worth getting all the innards out now so it's easy to build in later. This really is a beast, weighing nearly 17kg without any components in it, meaning it can be carried by one person, but you need to bend your knees when lifting it. The chassis also has removable panels on every side apart from the base, so it's good to just spend some time removing those as well. Most can be removed by hand, but some have screws and thumbscrews, so keep that screwdriver nearby, refer to the manual, and for the love of all things holy, make sure you keep everything together and separated from your other components.

2 CPU in, cooler prepped

Now, we turn our attention to the ASRock X670E Pro RS motherboard, which we take out and rest back on top of its box. We also grab the AMD Ryzen 9 7900X CPU, plus the NZXT Kraken Elite X73 RGB AIO. First, in shot 2a, the CPU goes into the socket with the label facing the right way up—just drop it on so it feels like it's in place and not moving, then lower the retention arm. Then, in shot 2b, we screw the standoffs into the motherboard ready for the cooler. We'd intended to use a different cooler for this build, but it was covered

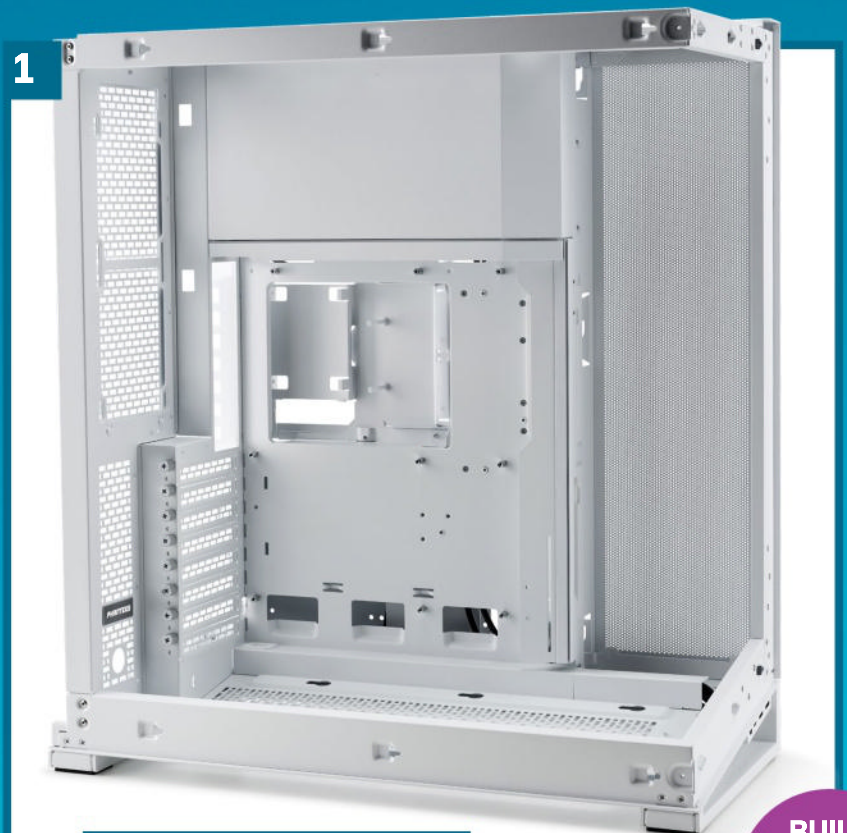

1

BUILD IT! STEP-BY-STEP GUIDE

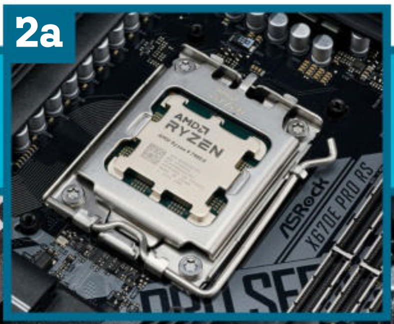

2a

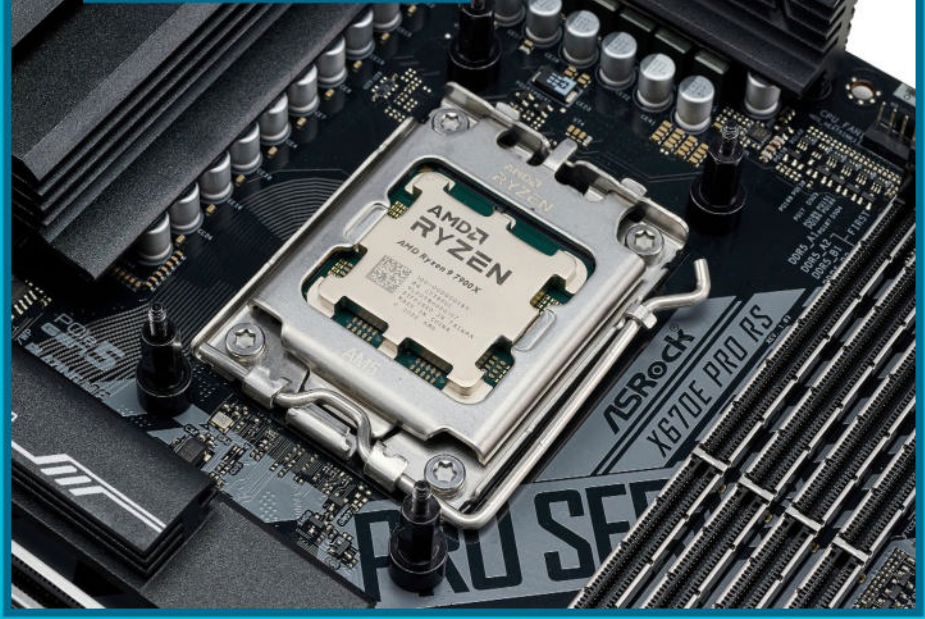

2b

3

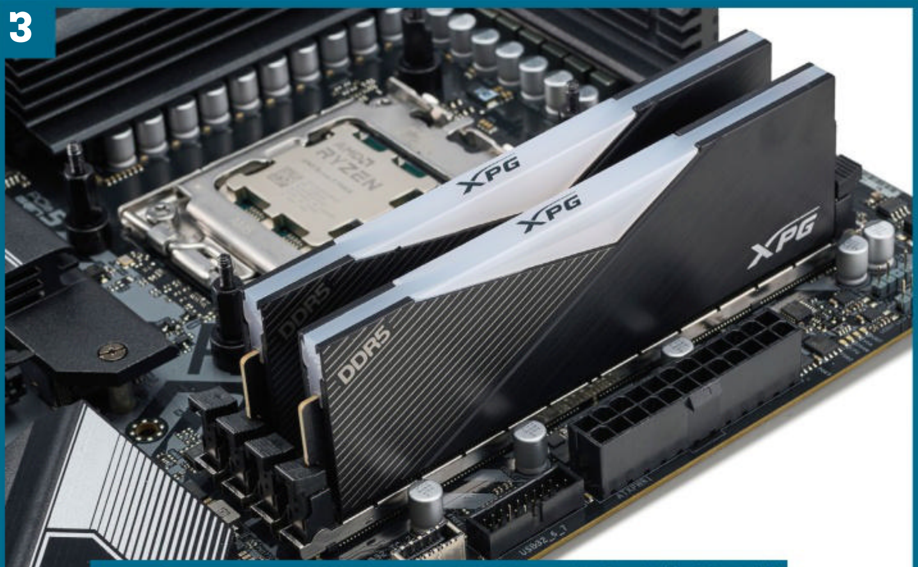

4a

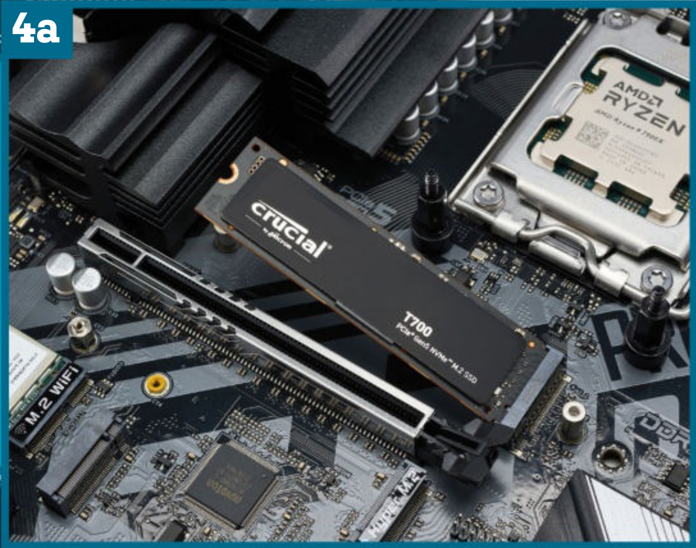

4b

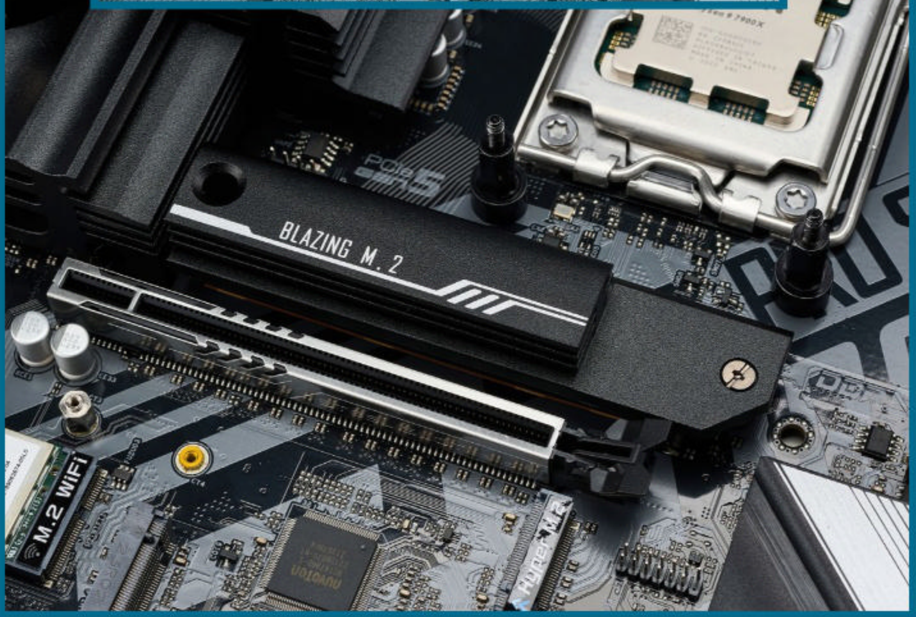

in dust and generally not in the best condition, so we repurposed the NZXT from a previous build. It's a super-simple installation with no backplate required—just the four screws you can see here.

3

Memory lights the corners of my mobo

We had a couple of sets of DDR5 RAM available to use in this system, and elected to use the ADATA XPG LANCER RGB 32 GB just because it was fast, and we'd used a lot of Corsair memory in the months previously. Most of you will know this, but just in case, you need to use the newer DDR5 memory with AMD's latest generation of processors, which will mean an upgrade is needed for many. However, don't be worried—DDR5 memory, like SSDs, has come down in price dramatically this year, so it doesn't command the premium it once did. ADATA's memory also boasts RGB lighting, which for this heavily lit system is kind of a necessity.

4

Gen5 Craziness

The most exciting part of the build so far (for me anyway, but I don't get out much) is the installation of the Gen5 PCIe storage from Crucial. This is the first time we've used a Gen5 SSD in a build, and we can't wait to see what it can do. While we don't know the benchmarks at this stage, we'd be surprised if we or anyone else will be able to tell the difference in most tasks between this and a Gen4 drive, which costs 50-60 percent less than this. So if you're looking to save any money at all on this build, we'd suggest starting here and going with a last-gen drive. But as we're *Maximum PC*, we can't wait to see what's being offered on the cutting edge of storage. This drive doesn't come with a heatsink, so use the one included on the motherboard, not forgetting to remove the plastic tab on the thermal pad first.

5

Our biggest fans

Another costly part of this build that you can skip or replace with a cheaper option are the 11 Phanteks D30 RGB fans we've used.

The generous folks at Phanteks were kind enough to send them all in with the case, which is large enough to actually fit one more in total along the front panel. This means we've gone for the luxurious option of fitting 11 RGB fans, including three on the cooler. We really didn't need to—NZXT's cooler comes with three RGB fans of its own, but we wanted a uniform look through the case. We also really wanted to try out Phanteks' new fan-linking tech, with plastic clips that make it really easy to lock fans together into a single unit, with only one power and one RGB cable needed for each row. The plastic clips provided then allow you to cover the screws for a super clean look throughout (see 5b).

6 Case, meet motherboard

If you want a better indication of just how big this case is, look at how it dwarfs our full-size ATX motherboard. I'm personally not too fussed about such things, but it'd be remiss if I didn't share that our photographer and designer both told me they'd have loved to have seen a white motherboard in this case for a cleaner design, so if a uniform look is important to you, please bear that in mind! Otherwise, this is a super-simple motherboard installation, again thanks to the case size—just angle it in so the back panel slots into place, then screw it in using eight of the screws provided in the Phanteks box. And with that, we can start connecting everything else that we need to the motherboard including power, cooling, and lighting.

7 Exhausting decisions

As a general rule, you want to position your radiator so that the top is above the highest point of the pump. Given the size of the case, that would be fine in the front or top of this case, but we went for the top, if only for a slightly cleaner look, bearing in mind that this build was also featuring on our cover. Looking to future upgrades, the top-mounting also completely takes up the space on that panel, while the front has space for one more fan, should we ever feel the need.

5a

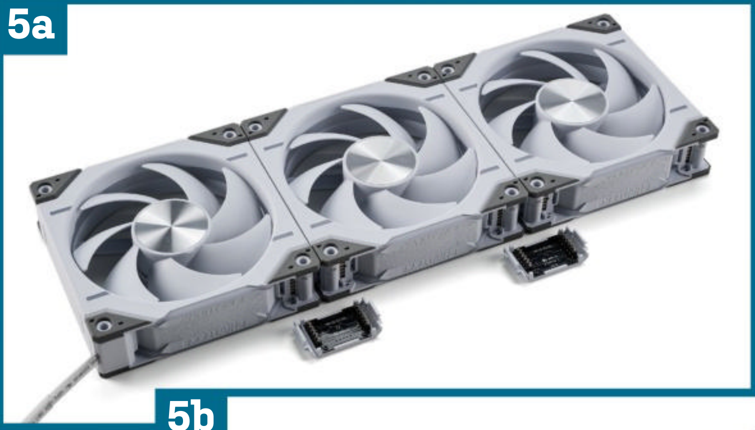

5b

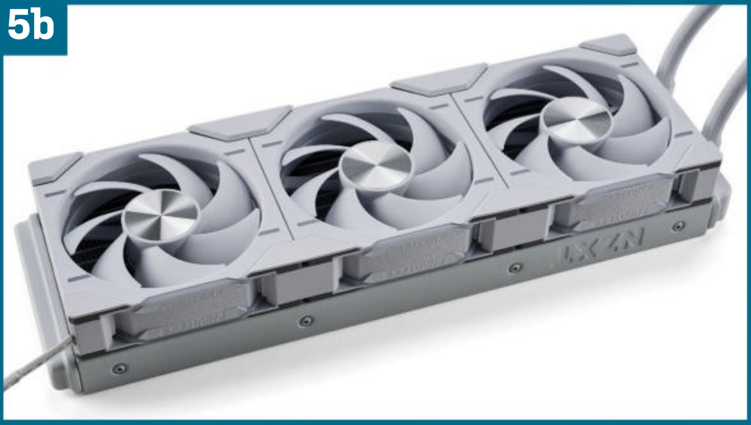

6

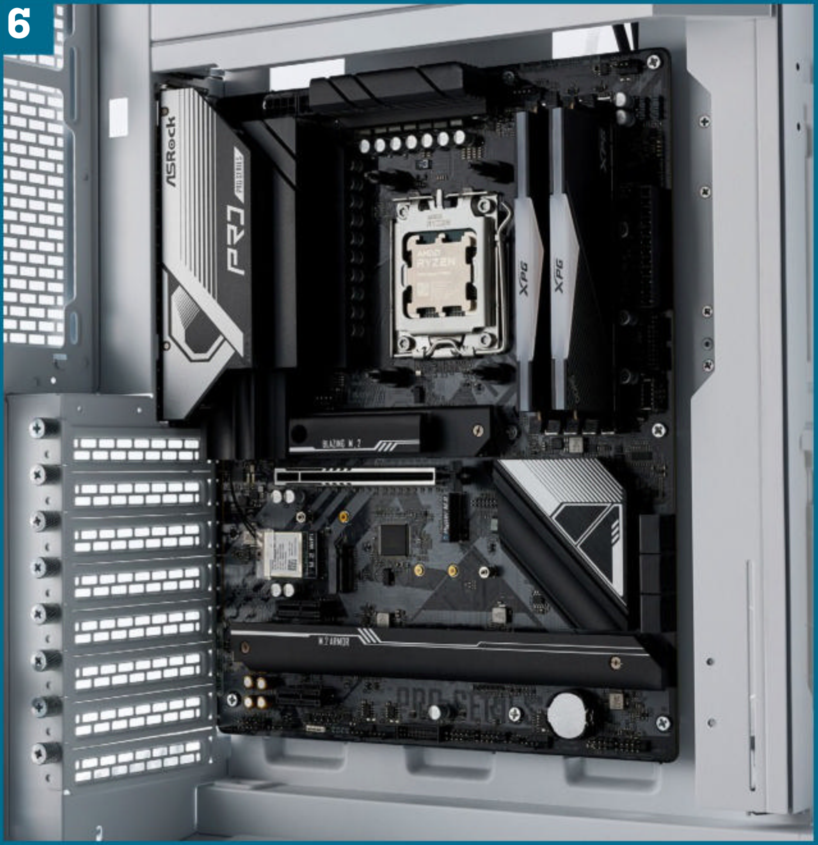

7a

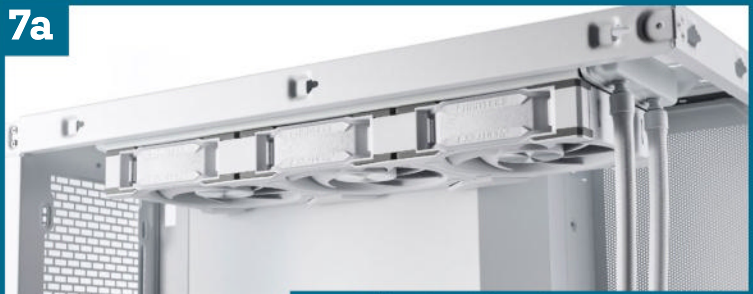

7b

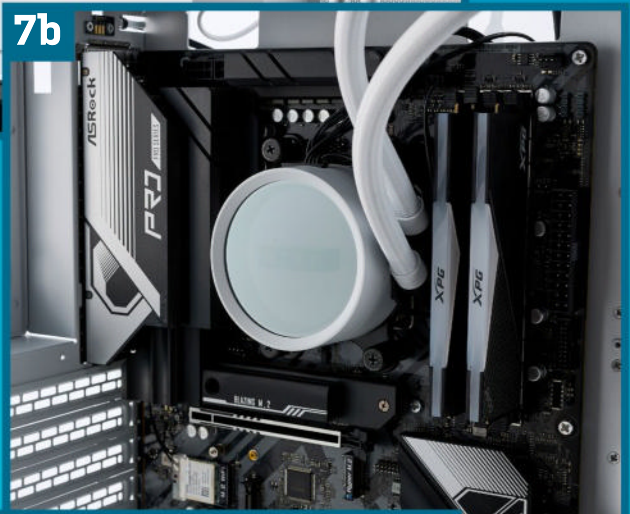

8a

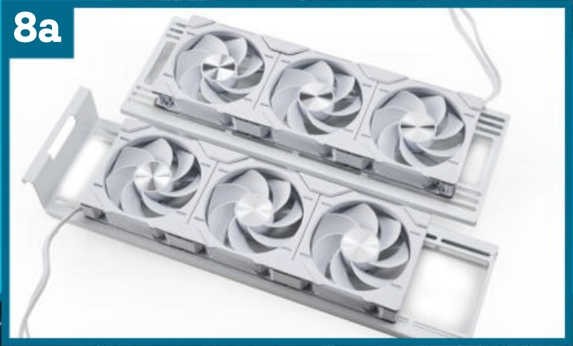

8b

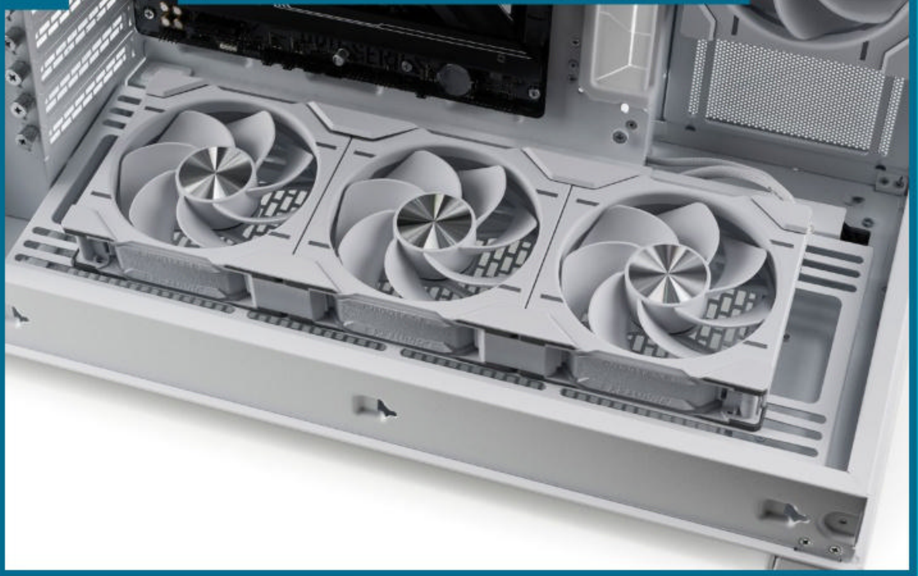

Installation of the rad is simple—just place it in the top of the case and screw it in from the outside using the screws provided with the radiator. Make sure the lighting and power cables run behind the motherboard to keep everything clean, and let the pump dangle while you're mounting in place. When you're ready to install the pump, apply a pea-sized amount of thermal paste (sold separately) onto the CPU, and screw it down into place.

8 Take me out

In our Phanteks case, the front and bottom fans mount onto removable plates, so you can completely discard them if you don't need them, or if you are using them, you can do the mounting outside of the case. The bottom panel (visible in 8b) has enough space for three 140mm fans to draw air in, although we only had 120mm fans supplied to us by Phanteks. The front panel has space for up to four 120mm fans, although as previously mentioned, we only had three D30 fans available. With 11 fans in this system, we shouldn't have anything to worry about when it comes to airflow, but completists may want to spring for the extra RGB fan if they want a fully tricked-out system. We love the way the fans snap together with the included connectors, as well as the caps that cover the screws for a clean look.

9 Light and airy

The final step for cooling is to install two 120mm exhaust fans at the back of the case, which is the maximum size and number you can fit at the rear. These clip together just as with the other D30 fans, and use one power and one RGB connector to run both, thanks to the daisy chain. Like the top fans, the rear ones screw directly into the case itself, but that's okay—there's plenty of room and clearance from the top of the case. We also like the number of cable routes in the case, as well as the matching white color of the cables, which means that you can keep everything neat and tidy as you do the build. Again, this is where we wish we'd plumped for a white motherboard to create a uniform look, but hey, at least we made this mistake so you don't have to.

10 Cable tidy

Now's a good time to start plugging in the cables from the case to the motherboard before we add the power cables and graphics card. As anyone who has built a PC before knows, some of these cables, particularly for audio and other connectors at the base of the motherboard, can be quite fiddly and difficult to connect first time. You want to be feeding the cables through the case's holes, connecting them to the correct port on the motherboard (see manual), and then pulling the slack back through to keep things neat and tidy. There aren't a lot of cables to tie up here, and as we aren't using any 2.5 or 3.5-inch drives in the build, we don't need SATA cables either.

11 Fiddle toy

Now, it's time to start connecting the power cables so that they're ready to be linked up to the PSU. The most fiddly is the CPU power cable, right up in the top left corner of the motherboard. However, thanks to this roomy case, it isn't as tricky as it has been in more compact builds before. Our photographer Olly has helped us use cable ties, which are included with the case in a little component box, to keep the rear fan cables tidy as well, by feeding one through the holes built into the case. The USB cable for the cooler, just to the right of the CPU power cables, isn't quite as tidy, but at least it's black, so it blends in against the motherboard.

12 Top-down approach

It's time to reveal our cable routing in the back of the case. The power supply in this case goes in the top of the machine, rather than at the bottom, as is more traditional. It's actually quite tight up there with all the power cables snaking their way to their components, particularly below the PSU. We suggest plugging all of the

9

ASRock
PRO SERIES
XPG

10

XPG
XPG
DDR5
DDR5

11

12

13

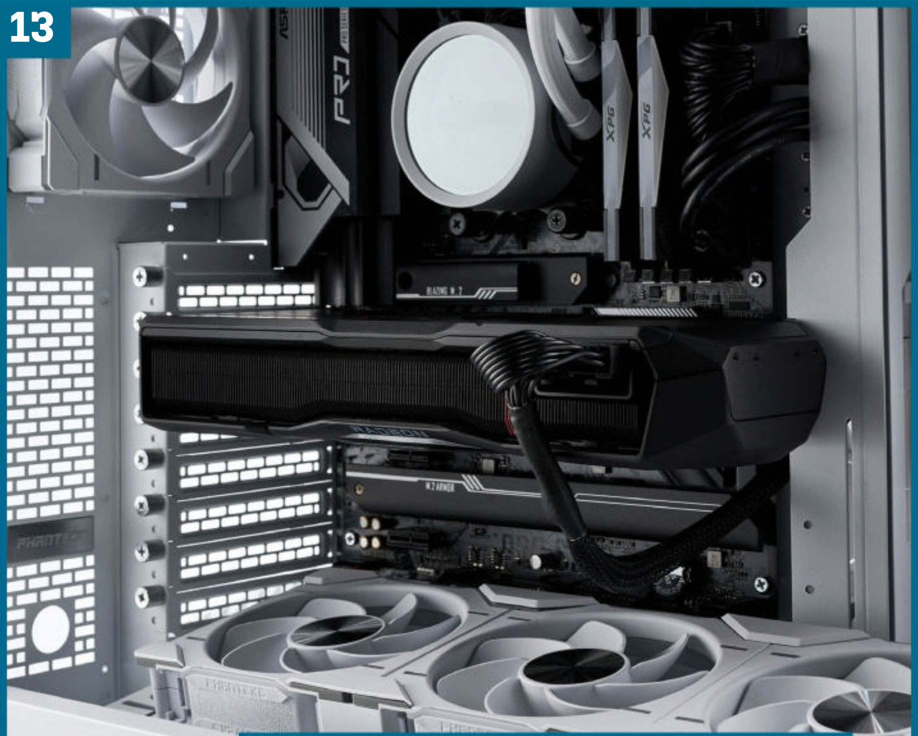

14

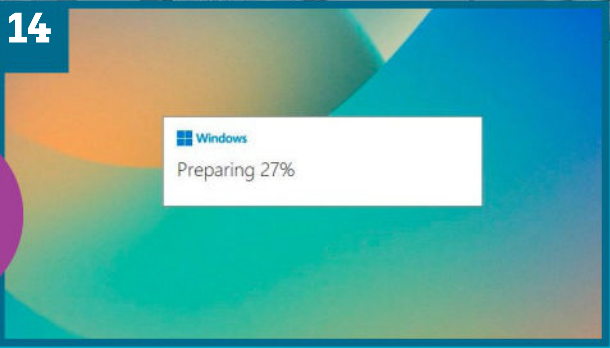

Preparing 27%

© MICROSOFT

BUILD IT!
STEP-BY-STEP GUIDE

power cables into the PSU while it's out of the case, then screwing it in once they're in place. The power cable and motherboard cables are routed downwards through plastic channels, with a rear door keeping everything neatly out of sight.

13 All-Out AMD

The final touch in this build is to plug the AMD RX 7900 XTX card into the motherboard, and connect up its power cables. Remove two of the rear shrouds on the case—you can use your fingers, but we found a screwdriver easier. Push it into the motherboard's top PCI slot until there's a satisfying click, and then tighten the screws back into the case to hold it in place. This card has two 8-pin connectors, which we'd recommend connecting to independent cables from your PSU—we could only find one when we built the machine, which was good enough for photography, but we made sure to use another when benchmarking for the best results. We didn't need to use one for the AMD card, but the Phanteks case also comes with a GPU support bracket for those heavier Nvidia cards, which is a nice little touch.

14 All's well that ASRocks well

We went through two broken power supplies on this build—one that made a whining noise when we booted, and another that did nothing at all. Luckily, we knew we had a working NZXT PSU and accompanying cables in another build, and thankfully they worked. Once we got the machine plugged into a monitor, keyboard, and mouse, we went to the ASRock website to download the latest firmware for the motherboard. We then searched online for "Create Windows Installation Media", and used a USB stick to create a bootable Windows installer. Windows installations are relatively quick and easy, as long as you have a decent internet connection. Once you're in Windows, be sure to download the latest ASRock motherboard and chipset drivers, as well as any important OS updates.

2023 high-end AMD PC

Great power at a reasonable price?

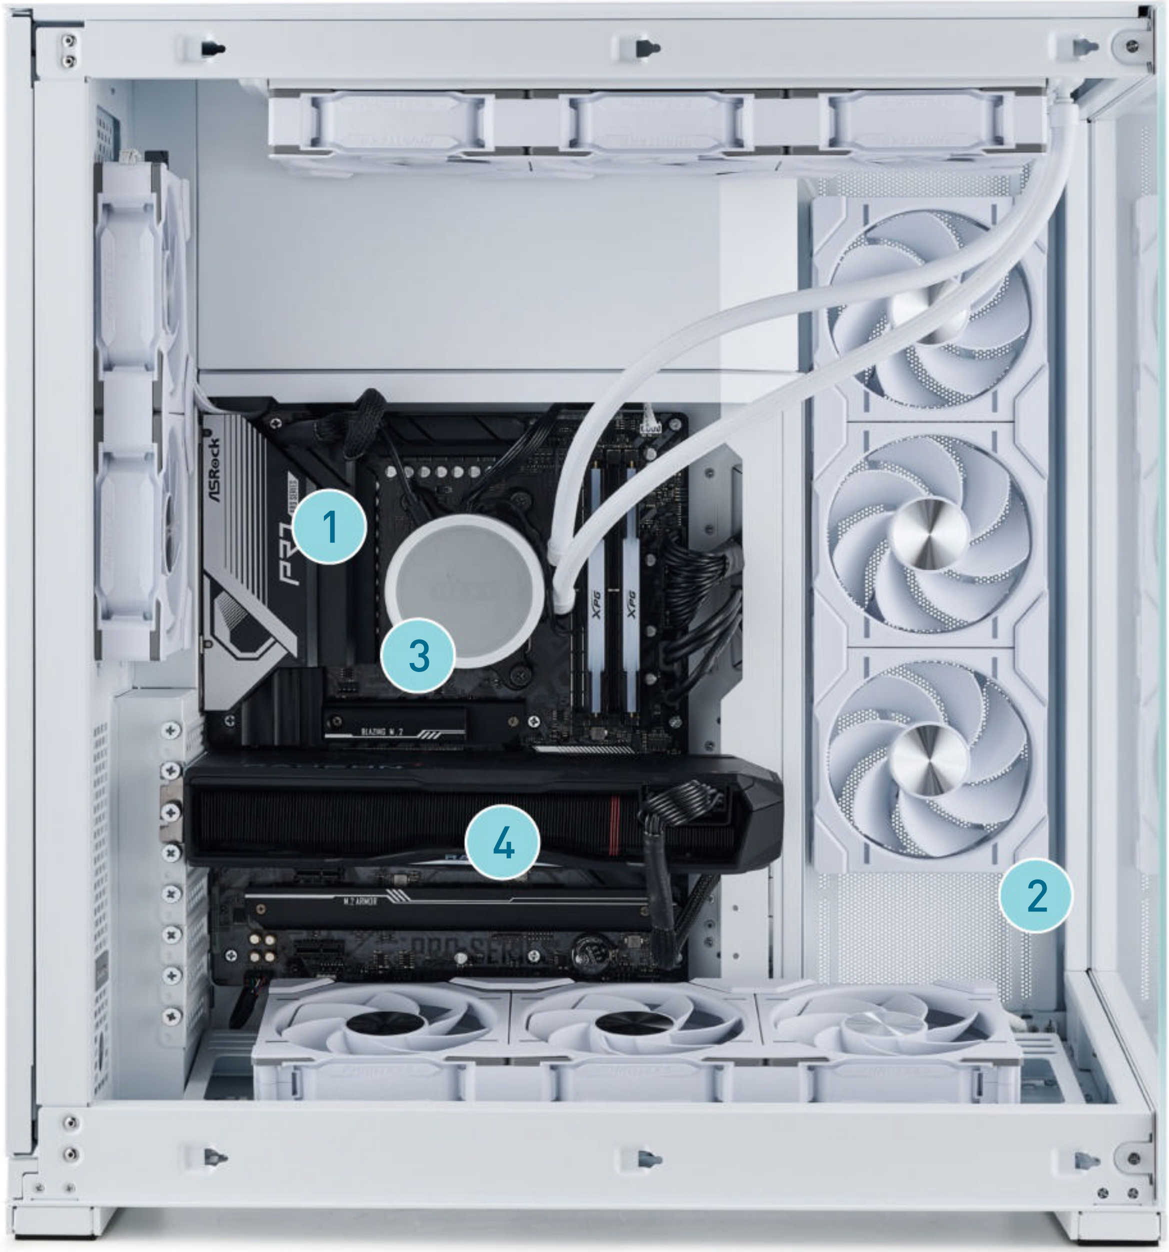

① The ASRock X670E Pro RS at $272 offers great value for money with all the newest features you could wish for. We just wish it also came in white.

② The completist in us wishes we had access to another D30 fan to fit in the front of this case, but to be honest, 11 total is plenty, and also quite costly.

③ The Ryzen 9 7900X is great for video rendering, 3D modeling, and virtualization, but if gaming is your primary objective, you might want an X3D chip.

④ This is our first build with AMD's RX 7900 XTX GPU. We love that you don't need power adaptors and support brackets like Nvidia's top-end cards.

AMD Gaming Bliss?

BEFORE WE TALK about performance, here are some notes on the build. We like Phanteks' NV7 case from a design and ease-of-build perspective—everything feels well-made and nice to look at. This is the sort of over-the-top PC that builders tend to love, and it attracted plenty of "oohs" and "ahhs" when we were shooting it in our studio. Equally insane is the weight. Sure, we could do with spending more time at the gym, but we struggled to lift this without help. This isn't a PC to build if you're attending lots of LAN parties.

When benchmarking, we compared it to our February 2023 "RTX 4080 Gaming Machine". We were tempted to pit it against the cheaper 7900 XT build from our April 2023 issue, but we think the 4080 and 7900 XTX are what most potential buyers will be comparing. With a $1,100 street price, Nvidia's RTX 4080 is still the more expensive card, which Nvidia would probably justify by pointing to its next-gen ray tracing and Super Sampling capabilities. However, at $200 less, can AMD's RX 7900 XTX make a case for itself when it comes to pure rasterization performance, plus the chance that it will one day catch up with its FSR 3 frame generation tech?

Let's answer the raw performance question first. In our suite of games, the AMD card exceeded the more expensive Nvidia card in 3DMark and non-ray traced versions of *Cyberpunk* and *Metro Exodus*. It was substantially poorer in the ray-traced version of *Cyberpunk*, a bit lower in *Metro* with RT, and slightly lower in *Total War*. Our recommendation would therefore be to save money and go for the AMD card if you don't care about ray tracing. If you want all of the pretty lighting effects, you're going to want the Nvidia 4080, or the 4070Ti at $800.

Other benchmark results are huge gains in CPU and SSD performance over our benchmark machine from February. The 50 percent performance improvement in multi-core performance between the $350 Ryzen 7 7700X and $450 Ryzen 9 7900X is substantial for anyone doing anything CPU-hungry, although we're surprised that performance didn't translate in the CPU-heavy *Total War* benchmark. The SSD performance gains were also substantial thanks to moving from PCIe 4.0 to 5.0, with much-improved read and write speeds. The caveat here is that we didn't feel any difference in game load times—*Cyberpunk* had the same load time as it usually does. If you need blisteringly fast transfers, then go for it, but for most, we'd still recommend 4.0 drives. While they will eventually be phased out, most people won't feel the benefit of PCIe 5.0 in daily use.

When it comes to aesthetics, we love the look of this build, although we did spend a lot of money achieving it. If you're more budget-conscious, which most mid-range buyers tend to be, then use cheaper fans, go for a PCIe 4.0 SSD, and find a less flashy cooler. That machine should then cost around $2,650—not bad for its insane polygon-pushing power, plus a CPU that's powerful without getting into the realms of silly money. While $2,650 isn't to be sniffed at, we think it offers good value for money—"insane performance without the insane price tag" is how we'd describe it.

BENCHMARKS

	ZERO-POINT	
Cinebench R23 Single-Core (Index)	1,962	2,016 (3%)
Cinebench R23 Multi-Core (Index)	18,502	27,796 (50%)
CrystalDisk QD32 Sequential Read (MB/s)	7,427	12,392 (67%)
CrystalDisk QD32 Sequential Write (MB/s)	6,649	11,809 (78%)
3DMark Fire Strike Ultra (Index)	31,648	34110 (8%)
Cyberpunk 2077 (fps)	73	78 (7%)
Cyberpunk 2077 RTX (fps)	49	20 (-59%)
Metro Exodus (fps)	91	101 (11%)
Metro Exodus RTX (fps)	63	60 (-5%)
Total War: Three Kingdoms (fps)	81	75 (-7%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of the "RTX 4080 Gaming Machine" from our February 2023 issue, featuring an AMD Ryzen 7 7700X, Nvidia Founders Edition GeForce RTX 4080 16GB GDDR6X, Asus ROG Strix B650-A Gaming WiFi ATX, 32GB of Corsair Vengeance DDR5-4800MHz, and 1TB SK Hynix Platinum P41 PCIe 4.0 SSD. All games tested at 4K 'Ultra' graphics presets with DLSS and V-sync turned off and XMP for RAM speed turned on. No manual CPU overclocking.

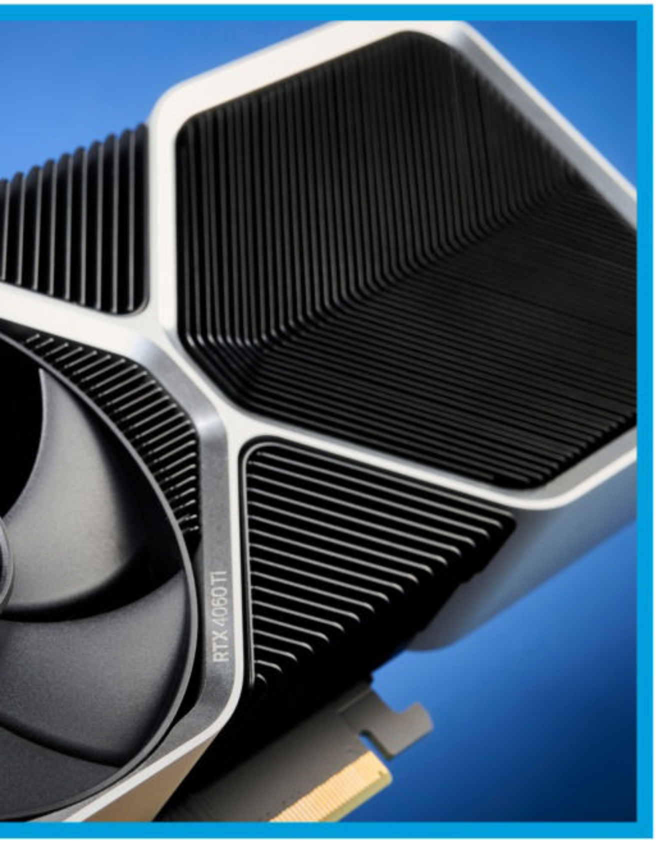

Cards like the RTX 4060 Ti feel increasingly under-equipped, with modern games pushing well beyond the 8GB VRAM barrier.

THE CHANGING STATE OF VRAM REQUIREMENTS

How much VRAM do you really need?

PC hardware trends upward over time. This is a well-established fact from the past four decades of home computers. Whatever you have now, you should expect your next PC upgrade to offer more: more performance, more memory, more storage, more features. Moore's Law suggested that we might get a doubling in capacity, performance, transistors, etc. every few years, but things have certainly slowed down over the past decade. What does that mean for our modern graphics cards?

The first consumer GPU to ship with 8GB of VRAM was AMD's Radeon R9 290X back in late 2014. That was the upgraded model; the original R9 290X (and 290) 'only' came with 4GB. Nvidia didn't join the 8GB club until late 2016 with the GTX 1080/1070 launch. Seven years later, we're at the point where 8GB feels very much like a budget-focused configuration.

And yet, rather than moving forward, the latest generation of graphics cards seem to be treading water or walking back VRAM configurations, particularly on the Nvidia side. RTX 3060 came with 12GB, yet here we are with the RTX 4060 and RTX 4060 Ti walking backward to 8GB configurations. Many gamers are now wondering how big of a problem that is right now, and what will happen going forward. But there's a lot more to VRAM than raw capacity, so let's dig in and look at everything going on with your GPU's memory subsystem.

HOW MUCH VRAM IS ENOUGH?

The question of how much memory your GPU needs and can effectively use has been around since the very first graphics cards. Way back in the day, Windows 3.1 could only run at 640x480 and 4-bit color on video cards with 256KiB of memory. Cards with 512KiB, meanwhile, allowed for 800x600 resolution with 8-bit color. That's because the frame buffer required enough memory to hold all the pixels at whatever color depth you wanted.

Modern PCs aren't just using a single frame buffer, but even the lowliest cards have plenty of memory to theoretically handle at least 4K resolutions with 32-bit color. A single 4K buffer only uses 32MiB, and we haven't had a new GPU ship with less than 2GB of VRAM since about 2013 (GeForce 700-series and Radeon 200-series). But VRAM stores things besides the various buffers now, and therein lies the concern with not having enough VRAM.

Games are the primary culprit, and besides multiple buffers for all the rendering—at least two frame buffers, a depth buffer, motion vectors, environment maps, and more—the textures used need to reside in VRAM. Well, technically they don't need to be in VRAM, but accessing textures from system memory can be more than an order of magnitude slower than having the textures in VRAM. If you're playing a game that tries to access 10GB of textures, buffers, and other data on the GPU and you only have 8GB, you get memory thrashing.

Anything not in VRAM gets pulled across the PCI Express bus, while something already in VRAM gets kicked out to make room. But then that evicted data might be needed, which means something else gets kicked out of VRAM and the other data gets pulled back in. It's a vicious cycle that can absolutely tank performance. That's because even the fastest PCIe connection for a graphics card only runs at up to 32 GB/s for PCIe 4.0 on an x16 interface, while a PCIe 3.0 x4 interface would only provide up to 4 GB/s. Even the lowly RX 6400 as an example has 208 GB/s of internal memory bandwidth.

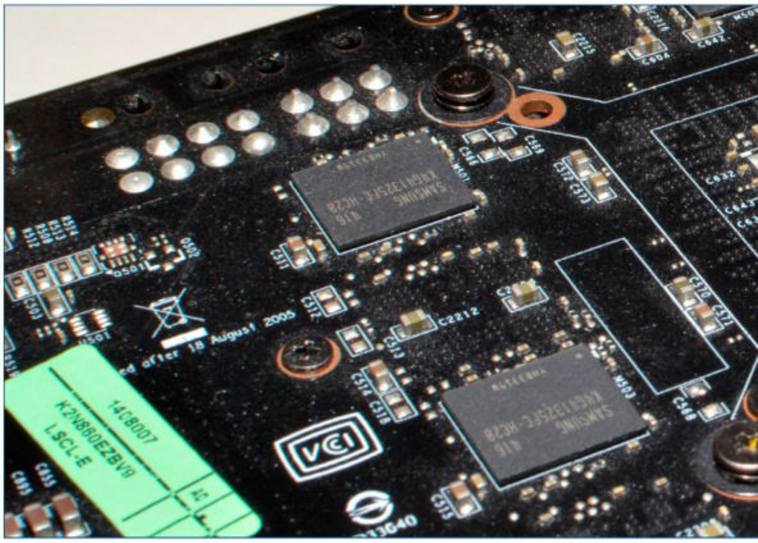

There are other scenarios where more VRAM can also be critical, large language models (LLMs) like ChatGPT being a prime example. GPT-3.5 reportedly needs to run in 800GB of memory, which requires multiple GPUs each storing big chunks of the LLM. Stable Diffusion XL as another example needs at least 6GB to manage 1024x1024 text to image generation. If your GPU doesn't have sufficient VRAM, the image generation will generally just fail with an out of memory error.

But let's go back to the subject of game textures for a moment, because it's not as clear-cut as you might think. Take a game with an HD texture pack, which usually means 2048x2048 textures,

The first consumer GPU to ship with 8GB of VRAM was AMD's Radeon R9 290X back in late 2014.

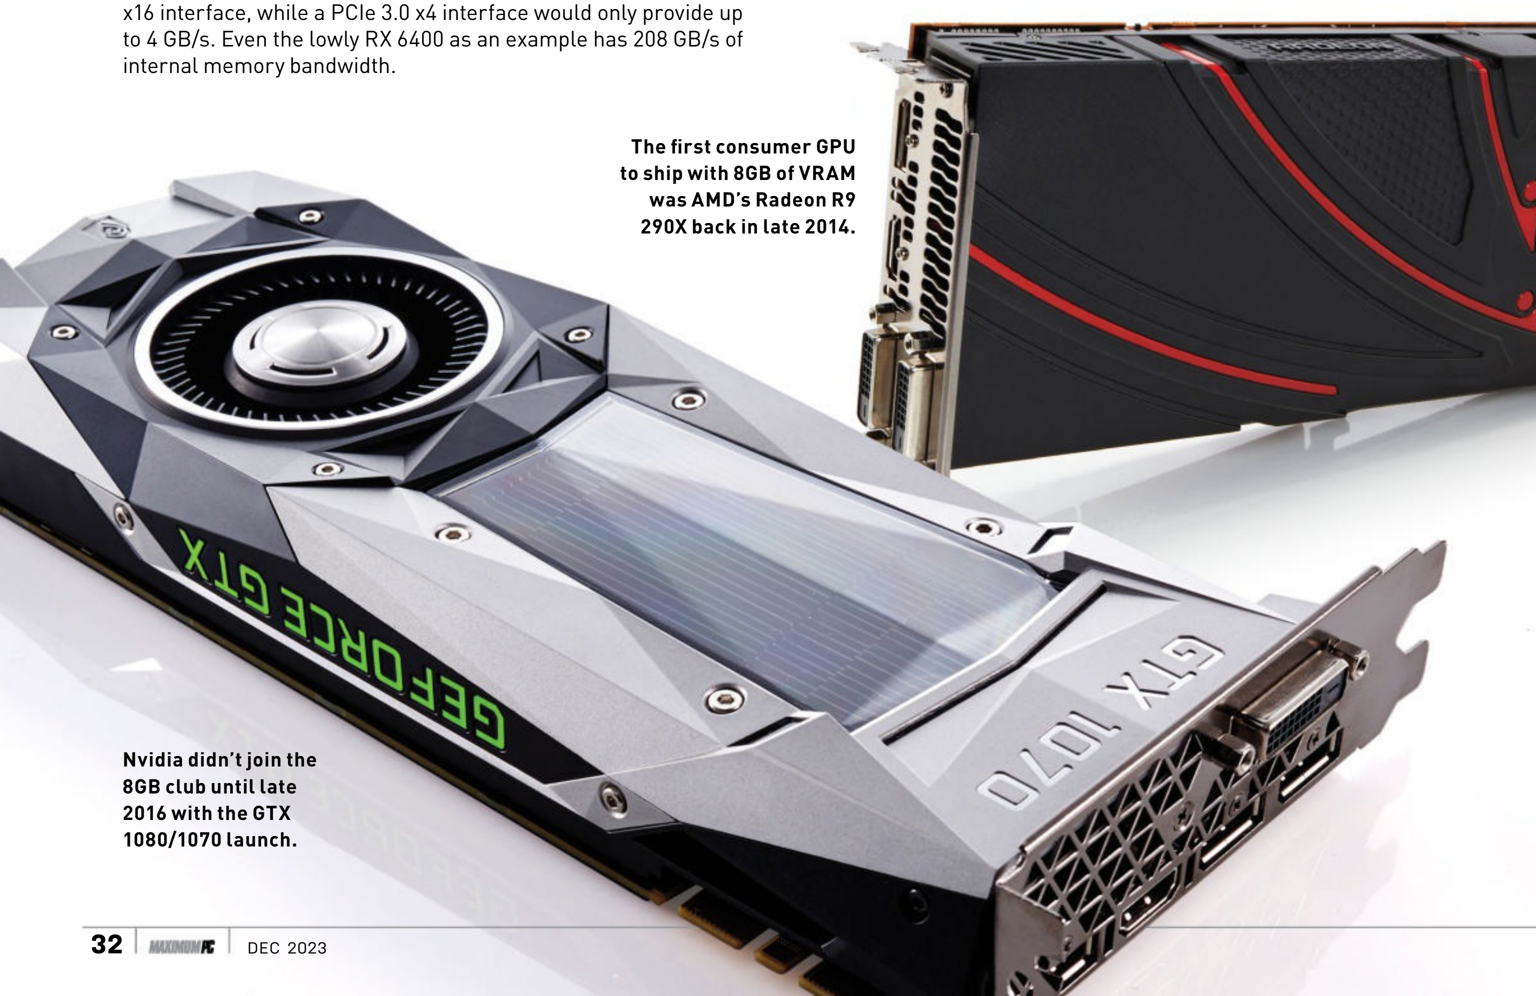

Nvidia didn't join the 8GB club until late 2016 with the GTX 1080/1070 launch.

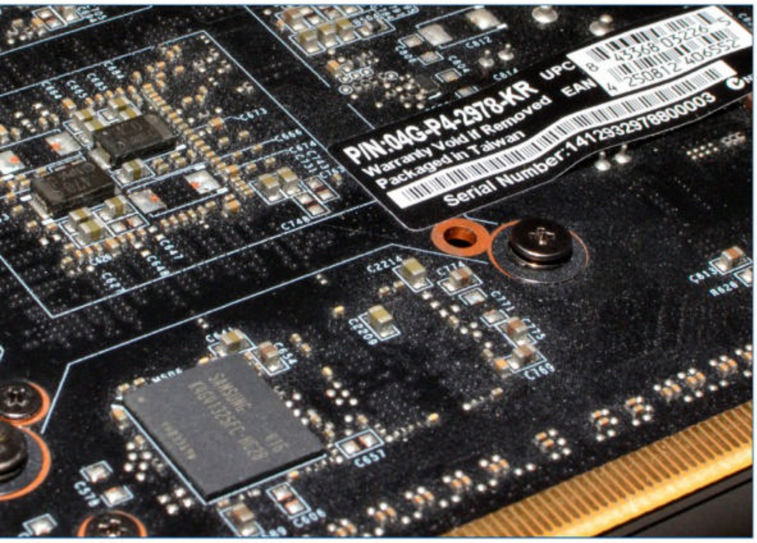

Putting memory on the back of the graphics card PCB used to be commonplace, like on this EVGA GTX 970, but now that's a "costly" manufacturing extra on cards like the RTX 4060 Ti 16GB.

sometimes even 4096x4096 resolution textures. Higher is better, right? The answer is yes, but it's a clear case of diminishing returns, and monitor resolution plays a significant role.

I've talked about this in the past (see Tech Talk in issue 218), but if you're running at 1920x1080, the difference between 1024x1024 textures and 2048x2048 textures often ends up being negligible. That's because the game engine will typically only use a mipmap for a texture that's one step above the size of the polygon (eg. a 512x512 mipmap will get used on a polygon that covers 400x215 pixels on the screen). Negligible in looks, at least, and there are exceptions to this.

Some games and engines do a poor job at texture memory management, and can end up with very low-resolution mipmaps on some surfaces and maximum resolution mipmaps on other surfaces. And then the high-resolution textures can shift around, so the end result looks very bad. But that's arguably more of a game engine problem than a case of not having sufficient VRAM—engines should be smart enough to manage memory use to get the best-quality output, and many have proven that this is absolutely possible.

MEASURING THE IMPACT OF VRAM

One solution to the above concern of running out of VRAM is to simply give every GPU more memory. It's not even that hard to do, though it does add to the cost — twice the memory chips costs twice as much. Just don't buy into the claims that it's somehow more expensive to put memory on the back of the PCB, as that used to be a common practice about ten years back. But can you have too much of a good thing?

The Nvidia RTX 4060 Ti 16GB card is the perfect candidate for investigation. It has twice the VRAM capacity as the standard 4060 Ti. At launch, it also carried a $100 price premium, though now it's down to a $50–$70 difference. This gives us a great option for seeing just how much memory capacity alone can help performance, keeping everything else basically static.

© NVIDIA

NEURAL TEXTURE COMPRESSION

We've been using compressed images and textures for decades now, but not all compression techniques are created equal. For photographs as an example, high-quality JPEG images can give visually similar results while using only 5–10 percent of the storage requirements. The problem is that JPEG compression isn't suitable for game textures. The quality might be good, but it's more complex to decompress the data, and it doesn't allow for random access to the stored data.

That last point is key for texture compression. Often, a game engine only needs part of a texture. With JPEG, it would need to decompress the entire texture before it could pick out the chunk it needs. To get around this, graphics APIs and GPUs support block compression algorithms. They're fast, and allow random access of the contents, but they don't compress data all that well.

Block Truncation Coding (BTC), used in virtually all games these days, hails back to the late 1990s. It offers varying compression ratios, but BTC high only provides 4:1 compression. That's better than nothing, but also explains why 8GB of VRAM can be limiting when games are trying to use thousands of 2048x2048 (2K) or 4096x4096 (4K) textures.

Nvidia is researching new algorithms that leverage AI networks to radically reduce texture memory requirements. Called NTC (Neural Texture Compression), Nvidia demonstrated storage of four times the resolution using 30 percent less memory. NTC could potentially reduce the texture memory footprint by around 80 percent while delivering similar quality, so an 8GB GPU with NTC would theoretically match a 40GB GPU using BTC.

NTC isn't ready for use in shipping games just yet, and it requires features like Nvidia's Tensor cores. It's also slower than BTC, and requires additional work from game artists. Still, it would be a significant improvement in texture compression. As one Nvidia spokesperson put it, NTC could fundamentally alter how future games and GPUs are made, which means that it's probably a decade away from being practical on most graphics cards.

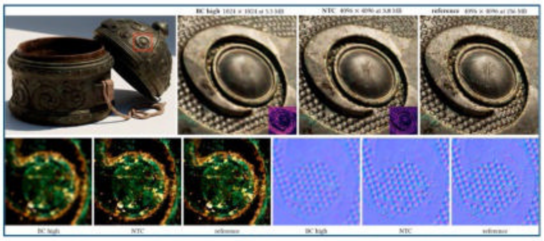
BC high 1024 × 1024 at 5.3 MB
NTC 4096 × 4096 at 3.8 MB
reference 4096 × 4096 at 256 MB
BC high
NTC
reference
BC high
NTC
reference

Nvidia is researching AI-powered options like Neural Texture Compression that could dramatically expand effective VRAM capacity.

16GB Vs 8GB

To measure the impact of VRAM, we've run benchmarks on the 8GB and 16GB variants of the RTX 4060 Ti and Arc A770. We've tested at 1080p, 1440p, and 4K, using high and mostly maxed-out settings—ultra, very high, or extreme, depending on the game.

		RTX 4060 Ti			Arc A770		
Game	**Setting**	**16GB**	**8GB**	**%Chg**	**16GB**	**8GB**	**%Chg**
Forza Horizon 5 (DX12)	1080p High	190.0	193.8	-2%	123.3	122.3	+1%
	1080p Extreme	115.6	117.6	-2%	81.2	74.4	+9%
	1440p High	155.9	158.1	-1%	105.3	104.3	+1%
	1440p Extreme	95.5	95.0	+1%	69.2	63.8	+8%
	4K High	100.2	101.7	-1%	73.2	72.6	+1%
	4K Extreme	63.6	59.4	+7%	48.2	43.9	+10%
Hogwarts Legacy (DX12)	1080p High	85.7	86.2	-1%	76.7	74.7	+3%
	1080p Ultra	72.7	71.8	+1%	58.5	56.2	+4%
	1440p High	66.0	67.5	-2%	61.7	59.6	+4%
	1440p Ultra	55.4	54.0	+3%	46.8	44.7	+5%
	4K High	39.6	39.1	+1%	35.1	32.9	+7%
	4K Ultra	32.6	30.6	+6%	28.9	26.1	+11%
Spider-Man: Miles Morales (DXR)	1080p High	97.3	97.0	0%	68.6	65.3	+5%
	1080p Very High	90.5	88.3	+2%	63.6	53.5	+19%
	1440p High	76.3	76	0%	56.5	50.5	+12%
	1440p Very High	68.5	59.0	16%	51.6	44.1	+17%
	4K High	45.4	41.7	+9%	36.2	30.9	+17%
	4K Very High	39.5	32.6	+21%	31.4	27.4	+15%
The Last of Us: Part I (DX12)	1080p High	86.1	85.6	+1%	46.9	43.4	+8%
	1080p Ultra	75.6	73.0	+4%	43.5	38.8	+12%
	1440p High	57.6	57.32	+1%	36.2	33.3	+9%
	1440p Ultra	50.5	48.4	+4%	33.1	29.5	+12%
	4K High	29.8	27.1	+10%	21.1	19.4	+9%
	4K Ultra	25.7	21.5	+20%	18.6	16.6	+12%

All testing performed with a Core i9-13900K on an MSI Z790 MEG ACE motherboard, with 32GB DDR5-6600 memory and a 4TB SSD. Tested cards: RTX 4060 Ti Founders Edition, Gigabyte RTX 4060 Ti 16GB, Arc A770 16GB Limited Edition, ASRock Arc A770 8GB.

(Clock speeds and cooling are slightly different, but the GPU core configuration remains the same.)

The RTX 4060 Ti isn't the only modern GPU to come with two different memory configurations. Intel's Arc A770 can also be found with 8GB or 16GB. It gets there by a slightly different path, however, as it has a 256-bit interface. The 8GB model has eight 1GB GDDR6 chips, while the 16GB model uses eight 2GB chips. The 16GB card also clocks the memory at 17.5 Gbps compared to 16 Gbps on the 8GB card, but we do get two different views of how performance can scale with more VRAM capacity.

Unfortunately—or perhaps fortunately—there's no third option using an AMD GPU. All of AMD's current generation chips only come with a single memory capacity configuration. The good news is that AMD generally gives you more VRAM at equivalent price points than Intel, and AMD GPUs that currently cost over $300 all come with at least 12GB of VRAM. Only the RX 7600 comes with 8GB, along with the RX 6600-class cards, and those all cost $250 or less and target 1080p gaming. The specs are more forgivable if the price is right.

To measure the impact of VRAM, we've run a collection of benchmarks on the 8GB and 16GB variants of the RTX 4060 Ti and Arc A770. We've tested at 1080p, 1440p, and 4K, using high and mostly maxed-out settings—ultra, very high, or extreme, depending on the game. We're intentionally using games that we know are more demanding on VRAM capacity, at settings that can exceed 8GB. But we also need to be realistic—neither of these GPUs are great picks for native 4K gaming, particularly in the most demanding games.

8GB VERSUS 16GB: FIGHT!

The four games we've tested don't represent every game out there, but the general trends are easy enough to spot. Contrary to what you might expect, the 16GB cards didn't demolish the 8GB cards. The 16GB cards did generally run better, but the biggest margins often came at questionable settings—settings that often dropped well below 60 fps.

COMPARING IMAGE QUALITY

We also need to consider not just performance, but how the games look with the chosen settings. That's hard to convey within the pages of a magazine, so we've uploaded images to the web to show the "high" versus "ultra" settings. Check out the full gallery at imgur.com/a/6RYVwM2.

The short summary: the visual impact of maxed-out settings doesn't often justify the VRAM cost compared to using high settings. There are aspects that may look better, but even that can be debated, as slightly nicer graphics but at lower levels of performance isn't a universally superior choice. Often, you can get nearly the same image quality by dropping the texture settings one notch and tweaking a few other options.

There are even settings like Chromatic Aberration and Motion Blur that reduce image fidelity in pursuit of "realism." There are guides for pretty much every demanding game on the internet, showing what the settings do and how much they impact performance. With a few judicious tweaks to your settings, you might improve performance by 50 percent or more without noticeably degrading the image quality. Nvidia and AMD even have recommended settings for your particular GPU that usually do a good of balancing image fidelity against performance.

The "ultra quality placebo effect" is very real, in other words. That goes double for things like ray tracing, where the improved lighting effects can often cut performance in half. Sure, *Cyberpunk 2077* in RT Overdrive mode with Nvidia's latest DLSS 3.5 Ray Reconstruction looks nice and runs decently on higher tier RTX cards like the 3080 and 4070. But it doesn't make the game more fun, and many people would prefer inaccurate lighting at 100 fps over full ray tracing at 40 fps.

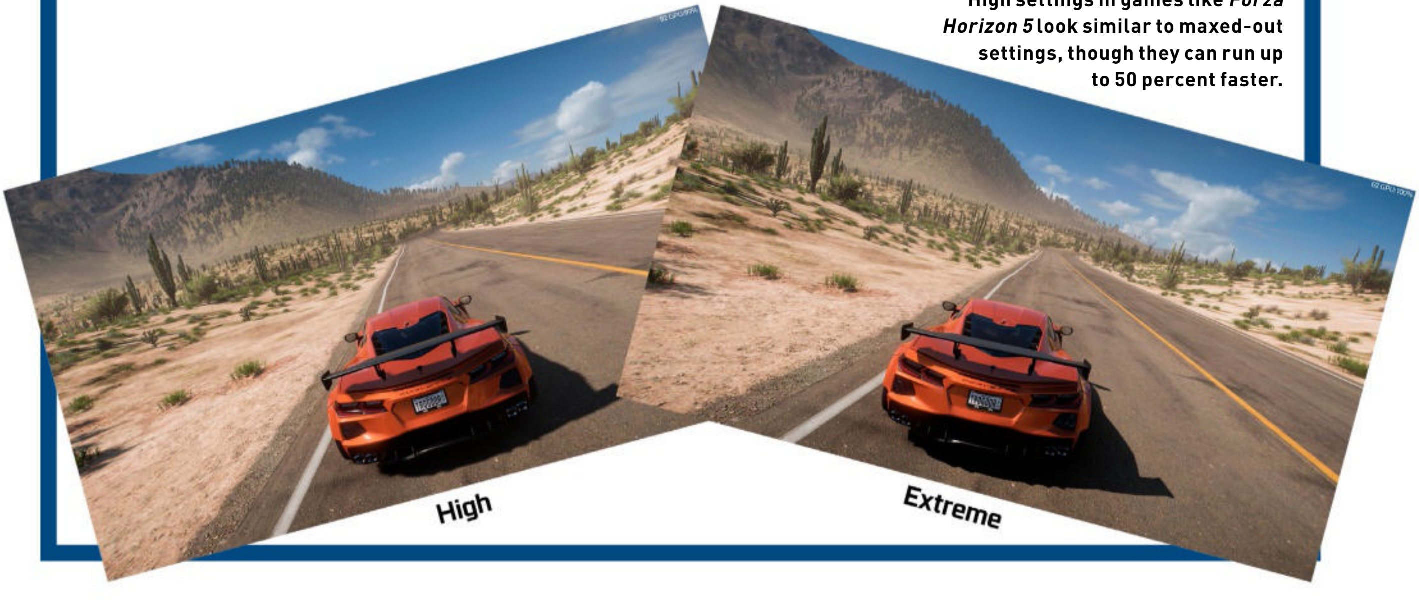

High settings in games like *Forza Horizon 5* look similar to maxed-out settings, though they can run up to 50 percent faster.

If you're only gaming at 1080p, 8GB of VRAM should still be sufficient, particularly if you don't try to max out every setting. For the RTX 4060 Ti 16GB, the biggest win at 1080p was in *The Last of Us*, where it was only four percent faster than the 8GB card. Intel's Arc A770 appeared to benefit more from the additional memory even at 1080p, but driver maturity likely plays a role, and the 16GB card also has nine percent more memory bandwidth. The 16GB A770 was anywhere from 1–8 percent faster than the 8GB at "high" settings, and 4–19 percent faster at "ultra" settings.

1440p shows some larger differences. *Spider-Man: Miles Morales*, for example, had the 16GB RTX 4060 Ti running 16 percent faster than the 8GB card at maxed out settings (including ray tracing). Again, Arc A770 either benefits from the extra bandwidth, or it just wants more VRAM capacity in general. Nearly every game ran measurably better on the 16GB A770, with around 10–20 percent higher performance. But *Hogwarts Legacy* and *Forza Horizon 5* at high settings only showed low single-digit percentage gains.

4K often represents the tipping point, but the 4060 Ti 8GB card still ran just as fast as the 16GB card in two of the games at high settings. *Forza* and *Hogwarts* at high settings were basically tied, while all four games using ultra settings ran 6–21 percent faster. For the A770, only *Forza* had the 8GB tied with the 16GB model, with a 7–17 percent lead in the other 4K tests.

BALANCED IS BEST

The key takeaway is that having "only" 8GB of VRAM isn't the end of the world, and at the same time stuffing 16GB or more on every

IT'S NOT JUST ABOUT AVERAGE FPS

We need to talk about minimum FPS, which isn't shown in the performance table. Even while average FPS was relatively close on the 4060 Ti, the 8GB card exhibited far more stuttering. *Spider-Man* was particularly bad, with one percent lows that were 50–80 percent faster on the 16GB card. For the A770, the gap wasn't quite as large in *Spider-Man*, though it still favored the 16GB card by 30–50 percent. One percent lows were also 15–20 percent higher in *Forza* using the extreme preset, and *The Last of Us* had minimum fps that were 10–35 percent higher across all tested settings and resolutions on the 16GB card.

We've noted many times over the past year that Intel's drivers aren't always up to snuff, and while our primary focus is on how memory capacity affects performance, it's also worth pointing out some of the discrepancies in performance between the Nvidia and Intel GPUs.

Across our normal test suite of nine rasterization games, all of which tend to be older and more established titles that likely received more attention from Intel, the RTX 4060 Ti outperforms the Arc A770 by about 30 percent on average. *Spider-Man* mostly lined up with that expectation, even with ray tracing, but *Hogwarts Legacy* had the A770 trailing by 20 percent or less, while *Forza* and especially *The Last of Us* gave Nvidia around a 50 percent and 70 percent lead respectively—and that's for the 16GB cards; the lead grew to around 60 percent and 80 percent if we only look at the 8GB models.

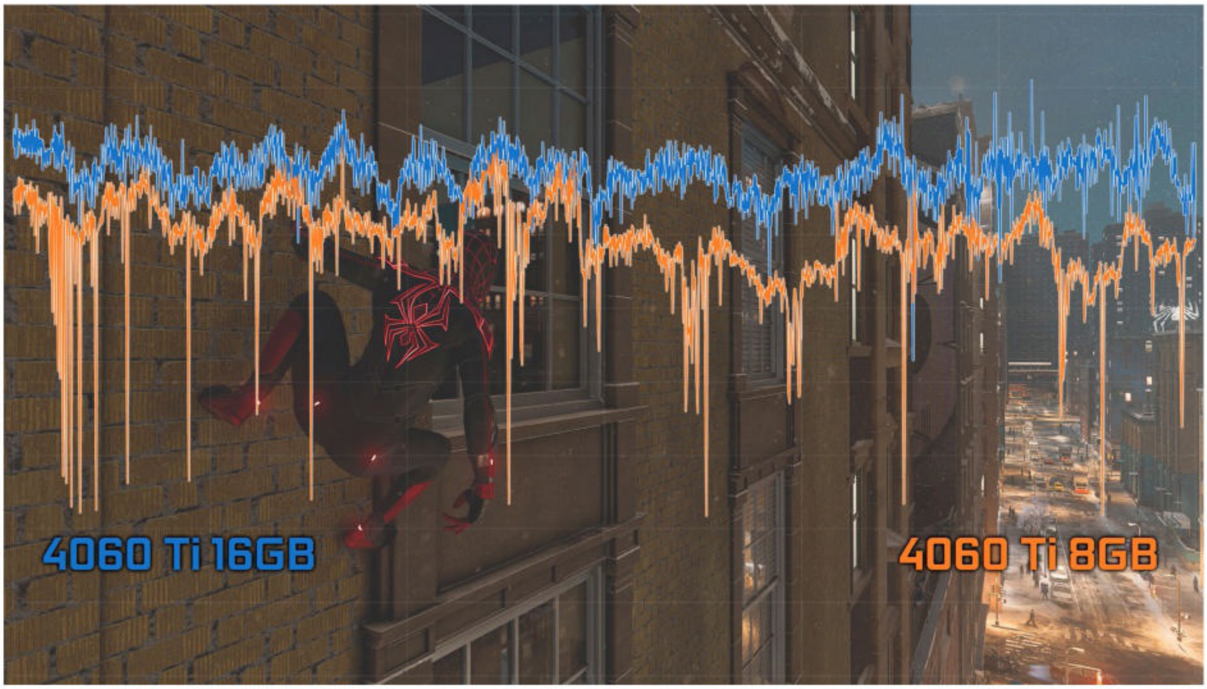

Here's a look at the FPS over time for the RTX 4060 Ti. All the large downward spikes on the 8GB card indicate stuttering.

©GEFORCE

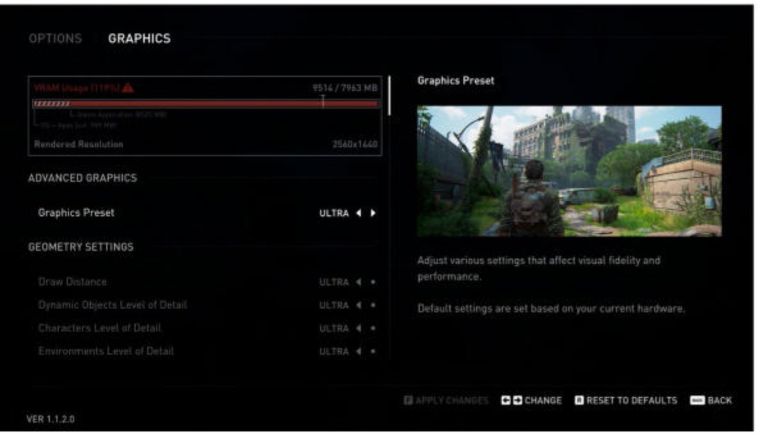

VRAM Usage (119%) 9514 / 7963 MB

Rendered Resolution 2560x1440

ADVANCED GRAPHICS

Graphics Preset ULTRA

GEOMETRY SETTINGS

Draw Distance ULTRA

Dynamic Objects Level of Detail ULTRA

Characters Level of Detail ULTRA

Environments Level of Detail ULTRA

Graphics Preset

Adjust various settings that affect visual fidelity and performance.

Default settings are set based on your current hardware.

***The Last of Us* gives a stern warning if you try to choose settings that exceed your GPU's VRAM capacity, like 1440p Ultra on an 8GB card.**

GPU going forward won't always help performance. In short, doubling the amount of memory per channel doesn't improve the GPU raw compute, nor does it improve bandwidth. The RTX 4060 Ti is still saddled with a 128-bit memory interface, whether it has four or eight 2GB chips spread across its four 32-bit channels.

More bandwidth on its own doesn't guarantee better performance either. Look at the Arc A770, where it has 78 percent (on the 8GB) and 94 percent (on the 16GB) more bandwidth than the RTX 4060 Ti, but still ends up trailing by often significant margins. If you really want better performance, you need more memory, more bandwidth, more compute, and perhaps an improved GPU architecture—something like an RTX 4070 or RX 7700 XT, for example.

Unfortunately, we can't simply pick and choose the specifications that we want. It would have been interesting to see an RTX 4060 Ti with a 192-bit interface and 12GB of memory, but Nvidia opted for a smaller GPU and fewer channels on the AD106 chip. The extra two memory channels would have increased die size quite a bit, as they would have also needed more L2 cache for the memory controllers. Since cache and external interfaces don't scale well with process shrinks, Nvidia instead opted for a smaller size and lower cost rather than a higher-performing part.

Ultimately, 16GB of VRAM today feels a bit like 8GB back in the days of the R9 290X. It can help, certainly, but it's far more a high-end option rather than something that even budget GPUs need. 8GB will still handle 1080p gaming, and it's a reasonable pairing for the lower-tier products. 12GB tends to be the sweet spot with balancing cost and capacity, and unfortunately Nvidia pushed that above the $500 mark this round, and even AMD wants $450 for the 7700 XT. But if you're looking at the RTX 3060 with 12GB and the RTX 4060 with 8GB, know that in every test at 1080p and 1440p that we've run, the 4060 wins, often by a wide margin—and that's without using DLSS 3 Frame Generation.

Our best advice is to try to get at least 12GB on your next graphics card upgrade, as that tends to be sufficient for just about any game running at up to 1440p. Future games will eventually push beyond that mark as well, but with the current-generation consoles sharing 16GB between CPU and GPU, most cross-platform games peak at around 10–12 GB of VRAM use. For better or worse, Moore's Law is slowing down a lot, but progress is still happening. It might be another ten years before we reach the point where 16GB is the minimum VRAM that we'd recommend, but that day will come, eventually.

© NVIDIA, THE LAST OF US

BUS WIDTHS, BANDWIDTH, AND CACHE

GPUs take years to go from conception to final product, and along the way, decisions are made regarding the features and specifications. The RTX 3060 used the GA106 processor, which has six 32-bit memory controllers. The RTX 3060 Ti used GA104 with eight memory controllers and came with eight 1GB GDDR6 chips. With the AD106 GPU used in the RTX 4060 Ti, Nvidia opted to stuff in a larger L2 cache while cutting down the memory controllers to just four.

That decision ultimately determines the possible memory configurations. For the base model, Nvidia uses four 2GB GDDR6 chips. The 16GB card puts four more chips on the back of the PCB to get to 16GB of total VRAM, but there are no in-between options that would allow for 12GB, like what we had on the RTX 3060 via six 2GB chips, and that's unfortunate.

It's not just about memory capacity, however. The 128-bit interface also reduces memory bandwidth. The fastest GDDR6 chips clock at 20 Gbps, and Nvidia uses 18 Gbps chips for the 4060 Ti. That yields raw bandwidth of 288 GB/s, compared to the 3060 Ti's 448 GB/s and the 3060's 360 GB/s. Reducing bandwidth compared to the previous generation would be terrible, but here's where the L2 cache comes into play.

Each memory controller is linked to an 8MB L2 cache, while the RTX 30-series used 512KB cache per memory controller. The 3060 Ti had 4MB of total L2 cache, while the 4060 Ti has eight times as much, even though it has half as many memory controllers. The larger cache means more data accesses get filled there instead of from the GDDR6 memory, which improves overall effective bandwidth.

The problem is that higher resolutions and larger textures reduce the effectiveness of the large L2 cache. That's why 1440p and especially 4K start to struggle more on the 4060 Ti relative to GPUs that have a 192-bit or 256-bit interface.

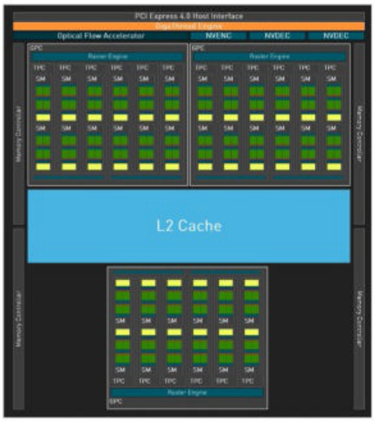

The block diagram for the AD106 GPU used in the RTX 4060 Ti shows that there's not much Nvidia can do to increase performance, as only two SMs are disabled.

GIFT IDEA!
Gifts that keep on giving

magazinesdirect.com

Over 100 brands to choose from

3 GREAT REASONS TO SHOP WITH US

- From tech and gaming titles to fashion and celebrity magazines there's something for everyone.
- Immerse yourself in our specialist one-off publications about your favorite hobby or interest.
- Exclusive Offer - Get 15% Off Guides and Specials - **MDMAG15**

VISIT OUR ONLINE STORE AT
www.magazinesdirect.com

✓ Convenient home delivery ✓ Save on shop price ✓ Never miss an issue ✓ 1.4m subscriptions sold

MAXIMUMPC

PREMIUM DIGITAL SUBSCRIPTION

PLUS receive a digital copy of these three special guides, and **100+ back issues**

PLUS
FREE
T-SHIRT
$500 AMD GPU REVIEWED PG.76
THE FUTURE OF CPUs PG.42
STEAM DECK UPGRADES PG.32
MAXIMUMPC
BEST VALUE GAMING PC
Get the most bang for buck in 2023
AMD RX 7800 XT and Ryzen 7 7800X3D
1440P FOR $2,000 PG.16
PLANT-BASED PCs PG.56
USE AI IN PHOTOSHOP PG.70
MAXIMUMPC
Linux & Open Source Annual 2023
retro GAMER ANNUAL
PC GAMER ANNUAL 2023

ORDER FORM COMPLETE ALL INFORMATION BELOW & RETURN WITH PAYMENT TO: MAXIMUM PC, PO BOX 2029, LANGHORNE, PA 19047-9957

SUBSCRIBE TODAY

☐ FOR FASTEST SERVICE
CALL 1-844-779-2822
AND MENTION CODE **B3DPREM**

☐ PAYMENT ENCLOSED
(CHECK PAYABLE TO MAXIMUM PC)

☐ 26 PREMIUM DIGITAL ISSUES
FOR $69.99

☐ 13 PREMIUM DIGITAL ISSUES
FOR $36.95

T-SHIRT SIZE
☐ M ☐ L ☐ XL ☐ XXL

BILL TO ADDRESS

NAME (PLEASE PRINT)

ADDRESS

CITY

STATE ZIP

E-MAIL

☐ Check here for news and offers from other Future brands.

*Offer is valid on US orders only, visit us at www.magazinesdirect.com or call 1-844-779-2822 for international pricing and other subscription options. Free gifts are fulfilled once payment is processed. T-shirt design may vary. 13 issues annually. **Offer valid until November 30, 2023.**

CLOUD COMPUTING DRIVING F1 FORWARD?

ORLEN
ALPHATAURI
HONDA
ALPHATAURI
ORLEN
PIRELLI
P ZERO
ORLEN
PIRELLI
P ZERO

Amazon's cloud computing resources helped F1 overhaul the design of the cars in 2022. *Barry Collins* examines whether it really made the sport a better spectacle

How do you make F1 more exciting? British media personality Danny Baker had the brilliant idea of making drivers do the first lap on foot, but that's unlikely to go down well with Lewis Hamilton and co. Instead, ahead of the 2022 season, the F1 authorities decided to use the power of cloud computing to make races more competitive.

Nobody wants to spend two hours of their Sunday afternoon watching cars file round a track in a procession—least of all the millions of fans drawn in by the deftly edited highlights shown in Netflix's *Drive to Survive*. With this in mind, Formula One decided to change the rules to enforce modifications to the design of the cars that would make it easier to overtake, inducing more of those moments that make commentator voices shoot up an octave.

F1's engineers and experts had a good idea about what modifications would result in more overtaking, but they needed the evidence to prove it. Testing is wildly expensive and terrible for the environment; wind tunnels are limited in size and scope, especially when you're measuring the impact on overtaking, a test that necessarily requires two cars. Instead, F1 turned to its partners at cloud computing giant AWS to see if software simulations could provide the answers.

Ahead of this summer's British Grand Prix, I met with Dr Neil Ashton, principal computational fluid dynamics architect at AWS, and Formula 1 pit lane legend and former Ferrari race engineer, Rob Smedley—now an F1 consultant working alongside AWS—to find out how F1 leaned on cloud computing to design the 2022 car, and whether it worked.

© ALAMY

LET THEM RACE

F1 often finds itself in a dilemma. The safety of the sport has improved vastly since the days when drivers such as Jackie Stewart had a 50/50 chance of dying over the course of their career. At the same time, fans aren't shy about reaching for the remote when races are boring. "Trying to get the potential of an overtake, or the actual overtake itself... creates the excitement that we all want," said Neil Ashton. "If we go back maybe a couple of years to the prior generation of racing cars, I think it was definitely the case that there was a lack of overtaking, and that was certainly reflected by the fans."

F1 teams design their own cars—it's part of what gives the well-funded teams such as Red Bull a competitive advantage over teams more often at the back of the grid. But teams must design those cars within the very tight confines of the F1 regulations, which stipulate how big certain parts must be, or put constraints on the design of different components. It was these regulations that the F1 management wanted to overhaul ahead of the 2022 season to encourage more overtaking.

Data has often been described as the new oil, and even in F1 the data is now much more valuable than the black stuff. When Rob Smedley joined F1 as a race engineer in the mid-1990s, "We still had abacuses", he joked.

"When I first came in, there were probably 20 sensors on the car," said Smedley. "Fast-forward not many years after—in fact, about five years—and there's probably something like 100 sensors on the car and like 1,000 or 1,500 data channels.

"If you look at today's cars, in test mode, they might have 500 sensors on the car, just to keep extracting more and more data."

Smedley said his generation of engineers pushed the use of data to "drive better decisions" for two reasons: "One, to design fundamentally faster racing cars and then, two, to optimize those racing cars for any given track or condition.

"We started to use much heavier deployment of data analytics—we were very early adopters of neural networks and subsequently machine learning. Given the complexity and how wide the problems are, it's just kind of gone on from there," he said.

With no shortage of data to fall back on, F1 bosses turned to the cloud to help them rewrite the rulebook.

TESTING TIMES

When it comes to testing F1 cars, there are several options, but none of them are cheap. Ashton said the "old-school way of doing it" is the physical test, where you drive a car around a circuit lots of times to collect data. "Formula 1 recognized that one, that was costing a lot of money, and two, it's not really environmentally friendly to be doing that sort of testing. So, physical testing from a Formula 1 point of view has long been reduced."

Then there are wind tunnels—dedicated facilities where big fans blow wind over

A familiar sight for F1 drivers: the back of Max Verstappen's car.

© GETTY IMAGES

the cars to test the aerodynamics of new designs. Ashton said NASA had one facility where you could fit a whole aircraft inside the wind tunnel, but "it requires the power of a small city". Although the F1 tunnels aren't that power-hungry, they're still expensive to construct and limited in space, which isn't ideal when you're trying to measure the impact on overtaking. "In this particular scenario, it was two cars following each other, which is definitely hard to fit in a wind tunnel in the first place," said Ashton.

What's more, when cars are out on the track, they're not racing in a straight line; they're barrelling round corners and chicanes, which are "very difficult to replicate in a tunnel that is usually a rectangle."

So, wind tunnels are expensive and can't measure half the things you want to check when evaluating the impact of new car designs. It would have to be simulated in software, but why go for the cloud instead of a dedicated supercomputer? "[Formula 1] did have access to computing power through an on-premise, more classical 'supercomputer', but that had a fixed capacity, which meant that simulations were taking about 60 hours," said Ashton.

With a supercomputer taking up to three days to run a simulation, and only having the capacity to run one simulation at a time, "It really did limit their capacity to go through many, many different designs," said Ashton. "Whereas the whole premise of cloud computing is you get access to this virtually unlimited capacity, where the limitation is really on the individual to produce the ideas rather than the computer."

Running the new car designs through a simulation on AWS's cloud infrastructure instead of the supercomputer reduced the time of each run from 60 hours down to 12. The cloud could handle multiple runs simultaneously. "That's the difference between getting something done on the same day versus two or three days," said Ashton. "You go into the office, you run something, you get it back by the end of the day."

The F1 engineering team originally planned to run 20-30 simulations a week, but was able to boost that to 80-90 tests with AWS's cloud infrastructure. With a limited window in which to rewrite the rulebook, so that teams had time to apply the design changes and run tests on their own cars, the cloud made a much bigger set of changes possible.

AS PURE AS THE WIND?

Running simulations on AWS might be cheaper and faster than on a wind tunnel, but is it as accurate? Neil Ashton claims it's close. "It depends on the specific simulation that's being done," he said. "It can be within a few per cent, so one or two percent difference. But it could be worse than that or better than that, depending on the specific technique that's used."

However, the more computing power that's thrown at the problem, the better the results. "There's a linear relationship between the accuracy of the simulation and the compute power required," Ashton added. "That's why so many companies are moving to the cloud, because they see that if they want to get their simulations to be more accurate, they need more compute. If you want to close that gap to 0.5 percent, it's increasingly going to be about more compute. That's the direction of travel across many different industries."

Aside from adding multiple cars to the scenarios, there are other things you can do in software simulations that you can't easily replicate in wind tunnels. "The additional advantage of simulation is the fact that you can start to include more environmental effects," said Ashton.

Real race conditions can't be replicated in a wind tunnel.

The Mercedes team's form tanked after the new rules were introduced.

"You could start to explore the effects of cross winds, the effect of the car being in front of you, as we said, the effects of temperature, the effects of rain."

Simulations also let testers model the fluid and structural motion of the car together, "So you can really start to get into the physics," said Ashton.

Finally, once all the simulations are run and results are collected, "That mathematical model can essentially go inside your racing game". That is, the car simulators F1 drivers train with. "When they give that test driver the steering wheel, that person will have pretty close to a feel of what the car would be like, so that when a team says we have a new part that we'd like to test, they will actually test it in the simulator first to see if it gives an advantage. If it gives an advantage in the simulator, then they may go and physically build it and send it to the track."

DID THE NEW RULES WORK?

Whether a F1 car can overtake the car in front of it (assuming there's enough space to pass) is largely dependent on downforce. Cars lose up to half of their downforce when racing closely behind a competitor, because of the turbulent air generated by the car in front, causing the chasing car to lose grip on the track.

The goal of the 2022 rule changes was to create a car with a much smaller "wake", reducing turbulence experienced by chasing drivers coming through dirty air. To that end, wheels were fitted with wake control devices; front wings were simplified to divert airflow off the front wheels, and the rear wing was redesigned to draw in air from the sides and lift the expelled air above the car behind. Wheel sizes were boosted from 13in to 18in, and fitted with low-profile tyres. All of this was simulated over six months, generating 550 million individual data points, before the rules were finalized and teams could get on with designing their own cars.

Did it work? A look at the season's final standings might lead you to conclude

GET READY TO BECOME A TRUE PYTHON EXPERT

Make Python work for you with tutorials on coding with Django, Flask, Pygame and even more useful third-party frameworks.

NEW
Everything you need to start coding with Python in Linux
The Python Book
Over 400 essential tips
ON SALE NOW
OVER 2 HOURS OF VIDEO TUTORIALS
Learn to use Python • Program games • Get creative with Pi
The Python Book Fourteenth Edition
100% UNOFFICIAL
ISBN 978-1-80308-208-0
FUTURE
FOURTEENTH EDITION

FUTURE

Ordering is easy. Go online at:
magazinesdirect.com
Or get it from selected supermarkets & newsagents

otherwise. Max Verstappen strolled to victory in the driver's championship, finishing almost 150 points clear of his nearest rival by winning 15 of the 22 races. The constructors' championship was equally one-stroke, with Verstappen's Red Bull team finishing 200 points clear of Ferrari.

The 2022 season was notable for the collapse in form of Mercedes. The 2021 championship was a nailbiter, settled in the final race of the season, when Verstappen overtook Mercedes' Lewis Hamilton on the final lap after a hugely controversial decision by the stewards to let only a small number of cars unlap themselves before the race was restarted. In 2022, however, Hamilton didn't win a single race and the team won only one.

Mercedes suffered badly from the rule changes. Early in the season, the completely redesigned car suffered terribly from "porpoising", a bouncing sensation that made the car difficult—even painful—to drive. Several teams suffered a similar effect, but the Mercedes car suffered the most, even porpoising along the high-speed straights.

But that doesn't mean the rule changes were a failure. Speaking after the fourth race of the 2022 season, several drivers said the rules had achieved the aim of making it easier to overtake. "It's 100 percent better," McLaren driver Lando Norris told the BBC. "It's what they intended to do and the reasons for the changes are working. I think you have been seeing better races this year than for many years." Verstappen and AlphaTauri's Pierre Gasly both agreed it had made it easier to battle the car in front.

Mercedes' woes might have been pinned on the rule changes, but there were other problems that haven't been resolved. Red Bull has won the first 12 races of the 2023 season, with Mercedes a distant second. The rules were tweaked again at the start of the 2023 season, and they will continue to be for years, as F1 looks to get the formula right. Hamilton's teammate, George Russell, remarked: "The sport took a really good turn for the better when these new cars were introduced [in 2022], but we need to take it to the next step now."

Neil Ashton thinks that's possible with ever-more cloud compute power coming online with each passing year. "If you could do more and more compute, you could design the cars more efficiently," he said. "That would create even more exciting cars, that would create even more exciting racing. I think we're in an inflection point now where it could really do something great for the sport." ⏻

BEHIND THE SCENES AT SILVERSTONE

When Formula 1 comes to town, a whole industry travels with it. Huge trucks arrive at each location, turning Transformers-style into temporary HQs. This is the paddock, where celebrities hang out and television hosts attempt to grab the gossip.

Tucked behind the paddock is F1's Event Technical Centre (ETC). In terms of data, this is ground zero of each race. Every single conversation between the drivers and their engineers, every byte of information tracked by the dozens of sensors on each car, each of the multiple 4K video streams that eventually appear on our screens.

While what you see on TV is topped and tailed by the hundreds of broadcasters around the world, they rely on a single source: F1. It has total control of the broadcasts, as if it was a TV production company. It also controls the technical services, from the timing that decides who wins and loses, to the telemetry that viewers see.

This process is managed by a comparatively small team with a few hundred employees. "We're not a massive organization," said Pete Samara, director of innovation and digital technology at Formula 1, speaking on the Friday (practice day) before this year's Silverstone race. "We're effectively an SME, and efficiency is really important."

At its most basic, that means buying hardware that will keep working ("that we don't need to switch on and off") and a system that allows staff to order whatever they need to their job. "We've got engineers, editors, broadcast graphics operators, CAD designers that are 3D-forming this building. So all types of user profiles."

For the past few years, F1 has chosen Lenovo as its hardware partner, and that partnership extends to its data center—or, strictly speaking, two data centers. We've already mentioned the mobile control center that pops up at each race, but that's merely a satellite compared to F1's Media and Technology Centre in Biggin Hill.

Processing the data captured at each event—around 500TB across a 90-minute race—needs more compute power than the hardware contained in the ETC, so it must be instantly beamed back to F1's expanded base in Biggin Hill. That data is then processed in real-time and transmitted to broadcasters.

F1 also has plans to enhance the quality of the data it provides. Rob Smedley said he'd spent some time with the soccer Premier League broadcaster. "You'll have your data scientists in the background saying 'Right, here are some really interesting stats', and then you'll get the pundits, who are actually engaging directly off-camera with the data science team... and they feed that directly to the audience.

"We need to get to that [in F1]," said Smedley. "There's technology we're trying to build with AWS that will get us to that. It will be more virtual, it will be more AI driven, because you've got 100 different broadcasters with 100 different commentators, all maybe at that point telling 100 disparate stories. So, in order to get in sync, then we've got to use the data, but we've also got to use technology."

The tech behind the scenes processes around 500TB of data at each F1 race.

© ALAMY

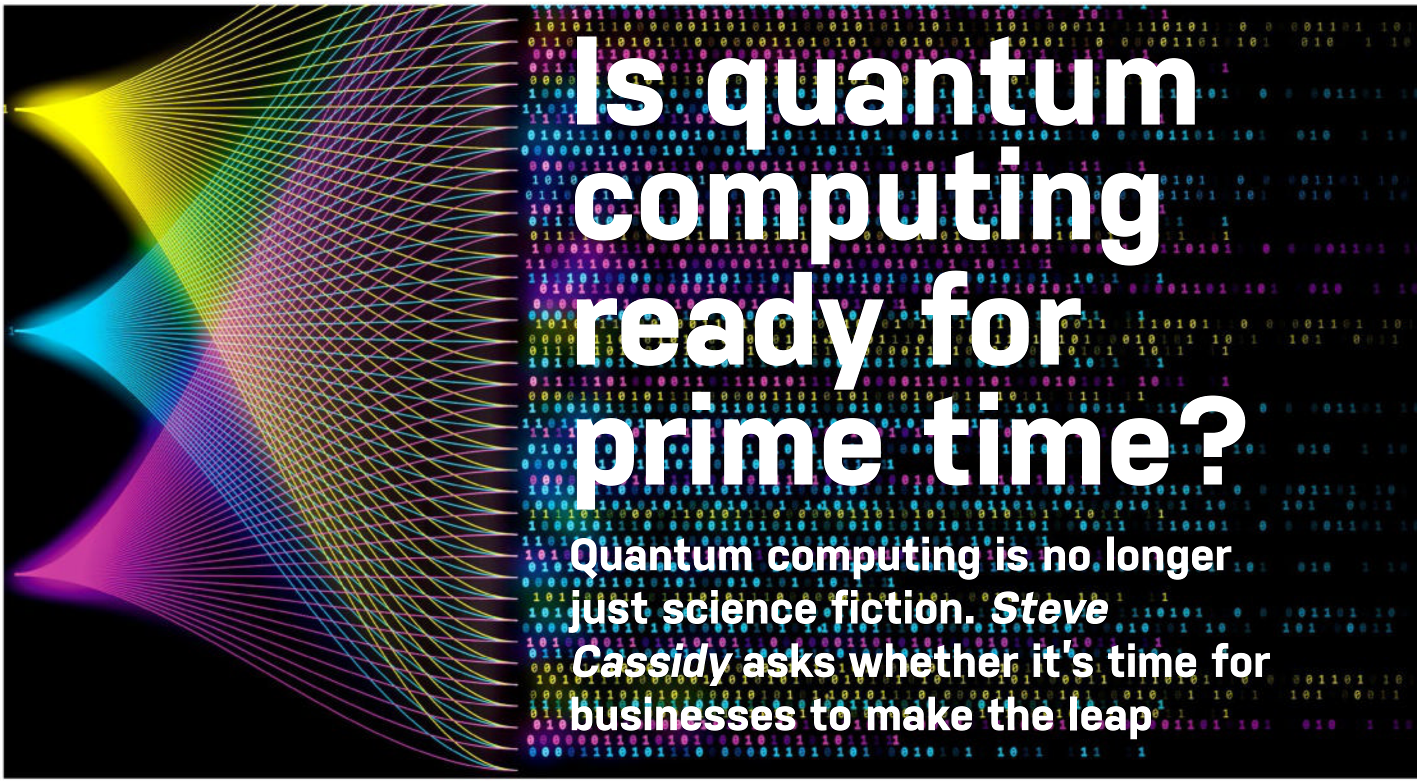

Is quantum computing ready for prime time?

Quantum computing is no longer just science fiction. *Steve Cassidy* asks whether it's time for businesses to make the leap

We're hearing more and more about quantum computing, and you might wonder what value it could have for your business. The answer is far from clear, because good, practical information about this new frontier is hard to come by.

I have to admit that I enjoy a nice information gap. Computing in general is stuffed full of such things, especially when new, specialist technologies collide with a less cautious and more gung-ho marketplace.

The general understanding of quantum computing (QC) doesn't go much further than the social media-driven perception that it represents the end of conventional computational barriers, and that leaps in performance are just around the corner for everything from tablets to data centers.

Of course, the more ambitious the perception, generally the less likely it is to represent the truth, and what we might call "quantum superstition" is no exception.

The potential of quantum

While the nascent technology might not transform every computing task, I've had a number of conversations in the past year that suggest it's time for businesses to start thinking about potential real-world quantum applications. That primarily means processes that involve performing analysis or computation on large data sets.

George Gesek, CTO at quantum specialist QMWare, gives the example of vehicle testing: when large manufacturers put their designs through the full range of product-testing procedures, they get back over 260 distinct statistics, settings, results, and predictions per vehicle. Some of those 260 data items might be single-number results, while others might be n-dimensional plots of fueling maps versus programmed power output management. Finding correlations and combinations of these values via conventional computing is intensive, repetitive, and slow; a quantum-based approach could do the same work in the blink of an eye.

Another discussion involved an American grocery store chain, and its habit of offering discount coupons. The question here was whether combinations of discount schemes might push the retailer into making a loss, and finding the answer was harder than you might think: coupons are a universe of fancifully ephemeral, limited data all to themselves. They can be local, global, time limited, product limited, modified by qualifying buyer, brand, age group… the number of permutations rises even more rapidly than the number of data points in the car-testing example, making it another ideal candidate for the particular strengths of QC.

As well as analyzing static data sets, quantum computing can be useful for modeling complex systems such as financial investment portfolios and the chaotic markets in which they operate. In chemical research, quantum analysis can simulate interactions between molecules more efficiently than any current computer system. Amid all the talk of universes, infinities, and criminal masterminds, there are business benefits to be reaped, especially for forward-thinking companies that embrace the potential of QC.

Quantum resources

Even though quantum could be the next computing revolution, nobody's blowing up your phone to try to sell you a quantum computer. In fact, it's doubtful whether on-

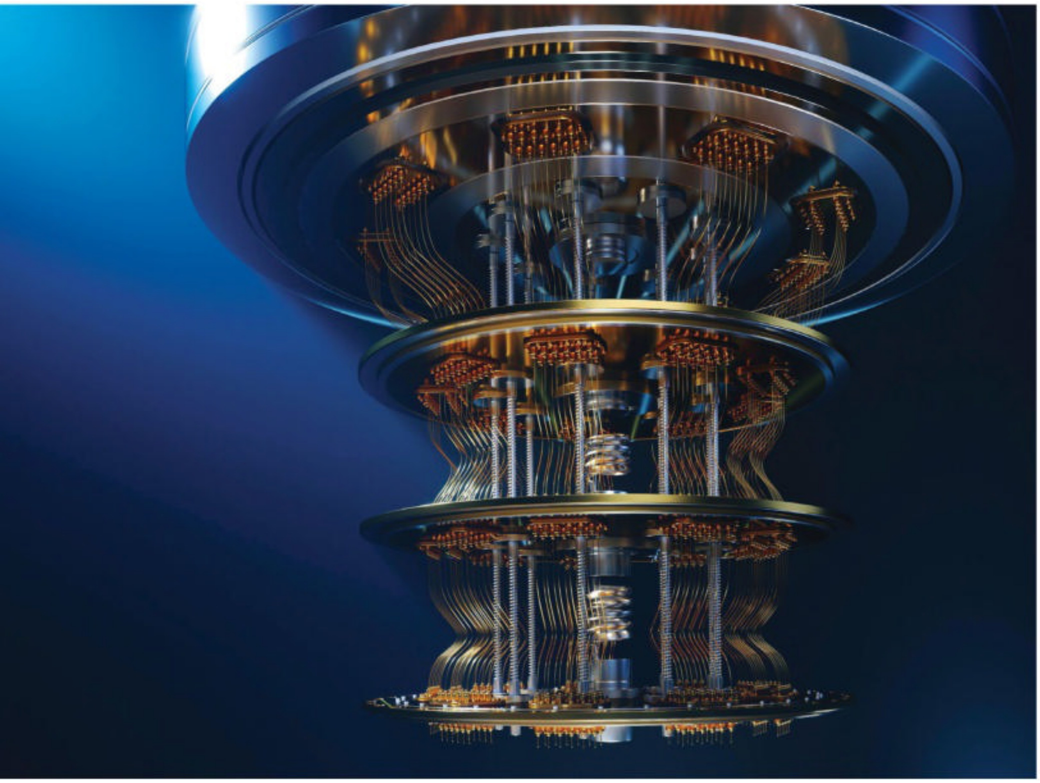

premises quantum computing will be a reality within two decades.

That's because while these machines have many similarities to regular electronic computers, the state of the technology is such that they're more expensive to build—I've seen $15 million per unit quoted. They're also harder to install and operate, thanks to the fact that separating the results of your QC computations from the background noise of the universe requires the core components to be kept a few degrees above absolute zero—around -270°C. So although the processing hardware might be no larger than your average pedestal server, the cooling apparatus can fill a room, and consume more energy than the rest of your business put together.

In practice, the only way for most humans to get access to a QC CPU is to lease access via a hybrid cloud architecture, from a supplier such as Amazon, IBM, Google, or Microsoft. This allows you to stay unfrozen at home, while your code runs in a data center in some ice-dripping cavern above the Arctic Circle. Existing and well-understood mechanisms keep the cloud copy and your on-premises versions in sync; the difference is the cloud version is reachable by the quantum CPU, while your stay-at-homes are not.

Preparing your workloads

So, you've signed up with a remote quantum provider—you're all set, right? Alas, no. There's another challenge that will likely be the dominant limiting factor in adoption for the foreseeable future, that being the brainwork of designing your application to run on a QPU.

This is not a trivial barrier. Quantum processors aren't just very cold Pentiums; their capabilities are tied to the underlying structure of reality (see "Quantum computing explained", right). There's no quick and easy way to convert an arbitrary workload into a quantum program. It can be a puzzle to determine whether your job spec is amenable to quantum processing at all.

I was intrigued by what my man from QMWare had to say, but when he started proposing an implementation of Turing's "universal simulator" concept—in which a punched paper tape can emulate any theoretical CPU you can imagine—I began to doubt whether this was within the reach of commercial possibility.

Quantum computers may look like sci-fi inventions, but they are now reality.

QUANTUM EXPLAINED

The word "quantum" is easy to misunderstand. Actually, it's almost universally misunderstood, because until the past few years, most of us had only encountered the term in science fiction. There the term has been thrown about quite wildly, perhaps to denote something that travels through time, or something that can be in two places at once. The latter image isn't wholly off the mark, but it doesn't do much to explain how the actual hardware operates.

The secret to quantum computing is that rather than regular bits that store ones and zeros, it works with "qubits" made out of subatomic particles. Using high-energy techniques that wouldn't be out of place at CERN, these registers can be placed into a quantum superposition of states, where they effectively represent both 0 and 1 at the same time.

By carefully manipulating a group of qubits, it's thus possible to carry out operations on hundreds, thousands, or billions of values at once, and almost instantaneously obtain results that might take far longer to find with a conventional computer. Eventually, the goal is to attain "quantum supremacy", whereby quantum computing is able to solve problems previously considered not computable within the lifetime of this universe.

While QC has revolutionary potential, not every type of operation can benefit from the amazing capabilities of qubits. Considerable specialist know-how is needed to identify suitable tasks, and prepare and present them in a way that can benefit from the technology.

QC AND AI

Perhaps the only buzzphrase bigger than QC right now is AI, so it was probably inevitable that they'd get mashed together. But this isn't mere marketing: QC has an important role to play, if not in general AI then in the subfield of machine learning.

I've said that the advantage of quantum computing kicks in when you're trying to analyse or optimize a large corpus of data points. As it happens, this is what a lot of machine-learning projects do: the process of training a virtual brain circuit to do the drudge work you have in store for it often entails some phases of very considerable compute burden. A quantum process called annealing can help out enormously by quickly processing a data set and finding optimal combinations of elements.

This might be an interesting source of pressure on the currently high cost of quantum resources, because ML is one of the most obvious general applications of the technology. It also promises to be a rich source of confusion, however: when one corner of AI rests on one application of QC, it's easy to imagine that this is all that quantum computing is good for—or that wider applications are just around the corner. In reality, the potential is both broader than this, and some distance away.

This is going to be a problem for as long as hybrid is the only realistic architecture for quantum access. As Gesek points out, the switch from on-premises to hosted, and hosted to elastic cloud, has been a smooth progression of demonstrable benefits. Moving to quantum cloud is going to be a harder sell: businesses don't often do things that are nerd-cool, unless it presents a clear route to making money.

The cost of quantum

That brings us to the next potential obstacle to widespread QC adoption: the price. It will depend on the tasks you're running, but you can be sure that we're not talking here about the everyday hosted services that run your website and a WordPress plugin or two.

It's not just the cost of the compute cycles, either. You may be used to the idea of a hybrid IT estate entailing extra costs to mix up local resources with the hosted cloud stuff, but to take advantage of QC, you need to go beyond a conventional hybrid outlook. You have one complete copy of your definitive business data lake held locally, and another where the supercooled chips can get at it, so you can run all of your apps, database services, third-party connectors, and the like within the purview of the QPU. This is the most expensive hybrid model possible.

The best way to achieve efficiency improvements may be by targeted simplification. Quantum computing is the ideal model for, say, analyzing a complete record of the behavior of all airmasses bordering the Atlantic for the past century—but that's probably overkill if you just want to know if it's going to be nice out today. As organizations refine the scope of the data they deliver to the hybrid space, it should be possible to reduce the required scale of quantum operations.

It will be clear by now that, while quantum capacity is no longer a mere sci-fi concept, there's a lot to work out before the technology can be called mainstream. Indeed, as in the early days of AI, the majority of supposedly "quantum" products and services popping up in the near future are likely to be nothing of the sort. But if you have an intractable modelling problem or a lot of data to pre-test for regulators, QC has the potential to provide immense gains—and you can start working on the code right now.

Quantum computers such as this one in Germany need room-filling cooling.

BAD QUANTUM

The one thing everybody "knows" about quantum computing is that since it can analyze any number of values simultaneously, it should crack the strongest passwords and encryption keys in the blink of an eye. There are also questions about whether quantum computing could blow apart virtual currencies such as Bitcoin, whose scarcity—and value—relies on its cryptographic security. Consequently, much of what you'll hear about QC in business is defensive. For example, Vodafone described its new business VPN service as "quantum-proof".

So far, we've seen no evidence of "quantum cracking". That's unsurprising: no black-hat hacker is going to have the funds and expertize to build and operate a quantum computer. Even if they were brazen enough to run such a workload on a leased platform, they'd need a brigade of PhDs for the code.

We might sooner see real-world attacks on Bitcoin, because cryptocurrencies are a threat to—and therefore a target for—nation-state hacking adventures. What angle the attack might come from is harder to predict. Governments may build their own labs, but it could be smarter to take advantage of the fact that cloud providers rarely know precisely who's making use of their resources, and to what end. That's always been true with traditional computing, and there's no reason to think it won't apply in the quantum age.

Join the conversation!

Discuss all things PC gaming with fellow readers at our official forums

forums.pcgamer.com

PC GAMER

THE POWER

Look beyond the mainstream options, and you'll find web browsers that offer many more features. *Barry Collins* examines three great alternatives

If you fired up Google Chrome today alongside the version that launched in 2008, you'd struggle to tell them apart. The address bar runs along the top, there's a row of bookmarks, a few buttons to the left, tabs at the top, and that's it. Has the way we use the web really not changed in 15 years? There are alternative browser makers who would disagree.

Google pioneered the fuss-free, feature-stripped browser in 2008, and the rest of its mainstream rivals followed. But there's a new undercurrent of "power browsers" designed to satisfy users who want more from their web browser: browsers with built-in email clients and feed readers; browsers that separate home and work life; browsers that don't endlessly distract you from the task at hand with notifications and alerts.

The web browser is the most used piece of software on most computers. Shouldn't it be better?

THE MOOD FOR CHANGE

Jon von Tetzchner has form when it comes to creating alternative browsers. In the early 1990s, he co-founded Opera, a browser that regularly won fulsome reviews for the breadth and originality of its feature set. In 2015, von Tetzchner did it again, founding the Vivaldi browser, and continuing the same ethos of delivering a feature-rich, privacy-friendly browser.

Von Tetzchner thinks that companies such as Google and Microsoft don't have to work too hard to make their browsers stand out because they have the inherent advantage of being able to bundle browsers with their operating systems. "For browser makers, the larger ones in particular, the point has always been, 'We don't need to differentiate; we just need to make it easy to switch. We have a distribution advantage... We want to make it as simple to change as possible.'"

There are other reasons why the big players don't pack lots of extras into their browsers. "If you don't have a lot of features, there's less that can go wrong," he said. "You can just focus on simplification, and then you can say, okay, if you want something more advanced than this, you have to go and get extensions."

© ARC,SIDEKICK,VIVALDI

BROWSERS

But in recent years, we've seen Microsoft begin to add features to Edge, and von Tetzchner believes that this is a consequence of consumers growing frustrated with the limitations of the browsers. "Microsoft has been trying to find a reason for people to use their browser, so they've been adding a little bit. Even Google has been adding features. I guess they don't want to get into a situation where they don't have any of the features and people actually start to move somewhere else."

Yet, even if Edge—with its newfound focus on AI assistance—and Chrome are beginning to pad out a little, they don't offer anywhere near the number of features and customizations you'll find in Vivaldi or emerging rivals such as Arc and Sidekick (you'll find full profiles of all three later in this feature). Yet all three of the power browser rivals are based on the same Chromium engine as the big two browsers.

Although the power browser makers must still be wary of feature bloat and its impact upon system performance, von Tetzchner believes adding features these days is less risky than it was in the past. "The speed of the network and the speed of the computer is less of an issue now than it used to be," he said. "Most computers are fairly advanced, and obviously now most of the browsers are using the same core, so there will only be limited difference in speed, unless you do something different."

ONE SIZE DOESN'T FIT ALL

For Dmitry Pushkarev, founder of the Sidekick browser, it's less about adding features to Chromium and more about taking things away. Pushkarev suffers from ADHD, and he wanted to design a browser that would help him work without the constant distractions. "Browsers have never been designed with work in mind," he told us. "In fact, they've been designed specifically to distract you, or send you towards the search engines. The way they make money today is basically they're being paid a fraction of search revenue. So, in fact, their incentive is very much to impair work and impair focus. Even though for most people it's just a minor nuisance, for myself and for maybe another 50 million people in the world, that's actually quite a bit of pain, because you really cannot get anything done."

Sidekick has a host of features designed to minimize distractions and keep the user focused on the task in hand, which Pushkarev claims benefits anyone working in a browser, not just people with ADHD. "The social cost of this pretty unbelievable—nearly a billion people work in a tool not designed for work," he said.

"You don't have any other industry like that, except for online work. Every other profession—take firefighters, restaurant chefs—they all have their own tools. We have a billion people who work online in tools designed to steal your attention. This is one of the big opportunities to make society better."

Want to see what an alternative browser could do for you? Our in-depth profiles of Arc, Sidekick, and Vivaldi might just convince you to switch.

© GETTY IMAGES

ARC

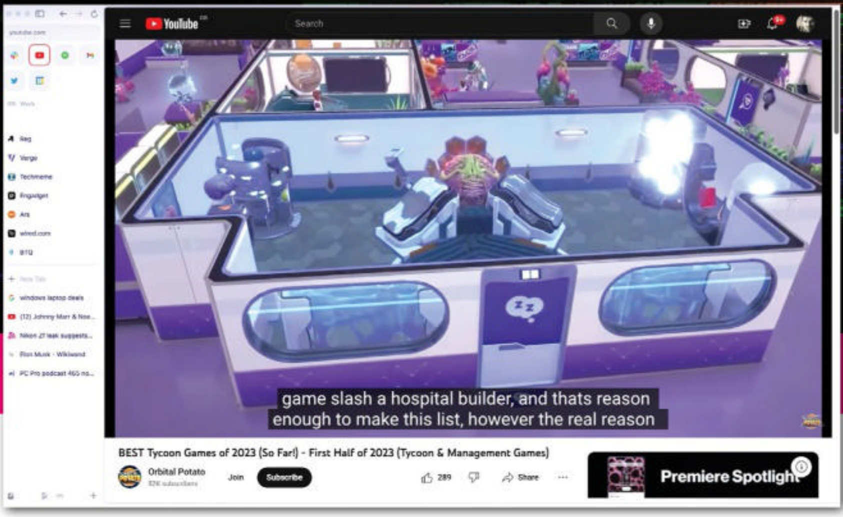

game slash a hospital builder, and thats reason
enough to make this list, however the real reason

BEST Tycoon Games of 2023 (So Far!) - First Half of 2023 (Tycoon & Management Games)

Arc packs a number of power features designed for enthusiasts.

Arc is the fashionable browser of the moment, the one that kids wielding MacBooks in Manhattan juice bars are using. I say MacBooks, because the browser is Mac only, although there are plans to release a Windows version.

Arc isn't just the browser to be seen with, however. It brings fresh ideas to the table, including power features that readers may well appreciate.

Arc is created by The Browser Company, a startup co-founded by Josh Miller and Hursh Agrawal, both former Facebook employees. However, it's another tech giant the pair have got in their sights. Visit the Arc homepage, and front and center is a pull quote from *The Verge*'s review of the browser, proclaiming: "Arc is the Chrome replacement I've been waiting for."

The company's mission statement on LinkedIn reinforces the message. "The world has changed in the past 15 years, but our web browsers look and behave pretty much the same. We think it's time to push the web browser forward again, which is why we built Arc—a browser that's not just faster, but also more personal, focused, creative... and maybe even more fun."

The one modern browsing problem that The Browser Company has identified is dozens of abandoned, disorganized tabs. The Arc browser is divided into Spaces, where tabs are organized by different themes or projects, and ordinary tabs that are left open for extended periods of time are automatically closed to reduce clutter.

"It's similar to your physical home—you have different rooms for different things, with different feels and different things in them," said co-founder Josh Miller in an interview with *VentureBeat* in 2021. "We want to give people that same level of control when it comes to designing their home on the internet. If my wife and I are trying to figure out our son's immunization requirements, that space should feel different and focus me differently than if I'm going down a rabbit hole on YouTube."

Miller claims that dividing the browser into Spaces helps to prevent attention drift. "I stay more focused when I don't have to context switch all the time," he said. "I have all these apps that I use, and I'm constantly flying around between them. It's much more distracting than people realize."

Arc looks very different to the big-name browsers, too. Most of the action takes place in the left-hand pane, leaving the right-hand side of the screen devoted to the web content. The address bar is squeezed into the top-left corner, just above a rather contained area for bookmarks. That said, a recent update allows users to "unhide" the toolbar, which puts the address bar back in its traditional place at the top of the browser window. Arc also makes it easy to tile tabs—you just drag them into position in the main browser window, much like split-screening in Windows.

Bookmarks work differently in Arc than in other browsers. When you click on a bookmark, its content fills the pane to the right, but it doesn't create its own browser tab. Certain sites and services also offer bookmark previews. For instance, you can hover your mouse over a Google Calendar bookmark and get a pop-up showing the remaining appointments in your diary for the rest of the day.

Then there are the more creative features in Arc, the ones that add the "fun" mentioned in the mission statement. Easels, for example, which give you a freeform web canvas on which you can place text, images, links, and shapes, or draw freehand.

© ARC

© GETTY IMAGES

ARC FACTS

Supported OSes: Mac, iOS (companion app), Windows coming in 2023
Browser engine: Chromium
Support for browser extensions: Yes (Chromium)
Founded: 2022

More recently, Arc added Boosts—an odd name for a feature that allows you to customize the design of certain websites every time you visit. Think the font on a particular site is too small? You can enlarge it with a Boost. Don't like the new Twitter "X" logo? Zap it and never see it again. (That noise you can here is a capricious billionaire imploding.) Boosts can be shared with other Arc users, too. One of the Boosts we found in the online Gallery removes the annoying Shorts from the YouTube homepage, for example.

Arc's biggest current weakness is that it doesn't really have a mobile version. There is an iOS "companion app" that effectively allows you to open pinned tabs in your Arc Spaces, but it's clunky, and not a great experience. Bookmarks are missing, there's no easy way to enter passwords on sites, and if you're using Arc on iPad, it appears at iPhone size, and has to be manually enlarged, which always looks awful.

The patchy mobile browser and current lack of a Windows version highlights one of the big problems for these startup browsers: getting stuff done without hundreds of dedicated developers isn't always easy.

SPACES

Spaces—similar to Vivaldi's Workspaces—help you organize browser tabs by activity. You might have a Space for home and another for work, for example, or different Spaces for different work projects. Each Space has its own unique set of pinned and regular tabs. Only favorites—the icons under the address bar in the top left of the window—remain in place, no matter which Space you open.

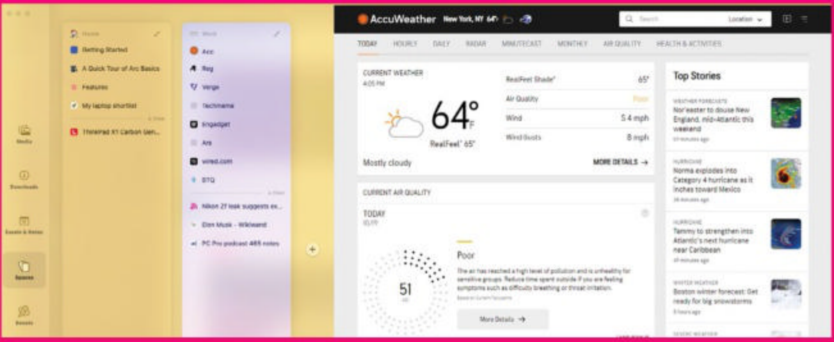

TAB HANDLING

Tabs are handled very differently in Arc. For starters, they live in the left-hand pane, not across the top of the window. Pinned tabs sit above the divider line, and are intended for those sites/services you're in and out of constantly, such as a web CMS. Regular tabs sit below the line, and are automatically expunged after a set time (the default is 12 hours) to prevent the Space becoming cluttered with open tabs.

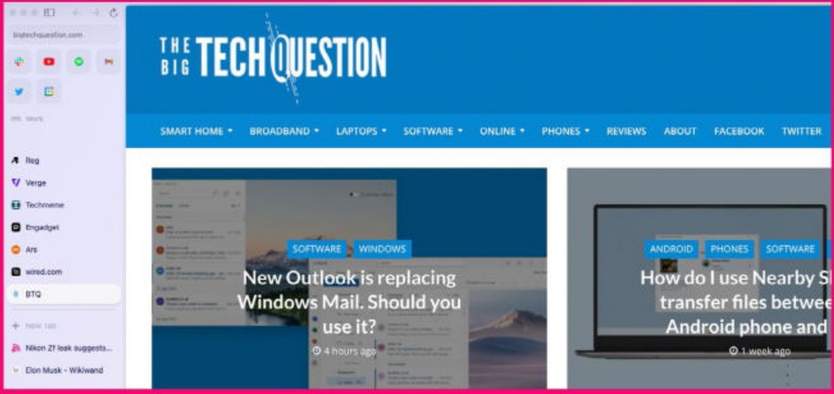

New Outlook is replacing Windows Mail. Should you use it?

4 hours ago

How do I use Nearby S transfer files betwee Android phone and

1 week ago

EASELS

Easels are a blank canvas, allowing you to plonk images and screenshots on a web page, annotate them with text, and draw arrows and other shapes. They're a great way to pull together ideas for a project. Easels can be shared and are viewable in any browser, but you'll need Arc to edit them. Arc also has a more conventional Notes facility.

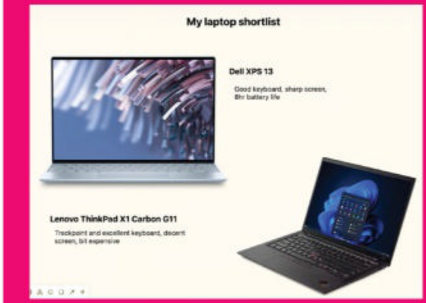

My laptop shortlist

Dell XPS 13

Good keyboard, sharp screen, 8hr battery life

Lenovo ThinkPad X1 Carbon G11

Trackpoint and excellent keyboard, decent screen, bit expensive

LITTLE ARC

This feature helps prevent your browser drowning under the weight of open tabs. Say you're reading an article and there's a link to a related item—you can right-click on it and open it in Little Arc, a small window overlaid on your current window. You can quickly check the info on the link and pick up where you left off on the original page.

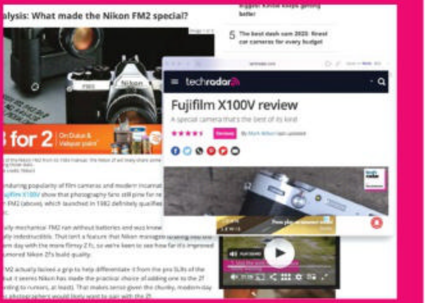

lysis: What made the Nikon FM2 special?

Fujifilm X100V review

SIDEKICK

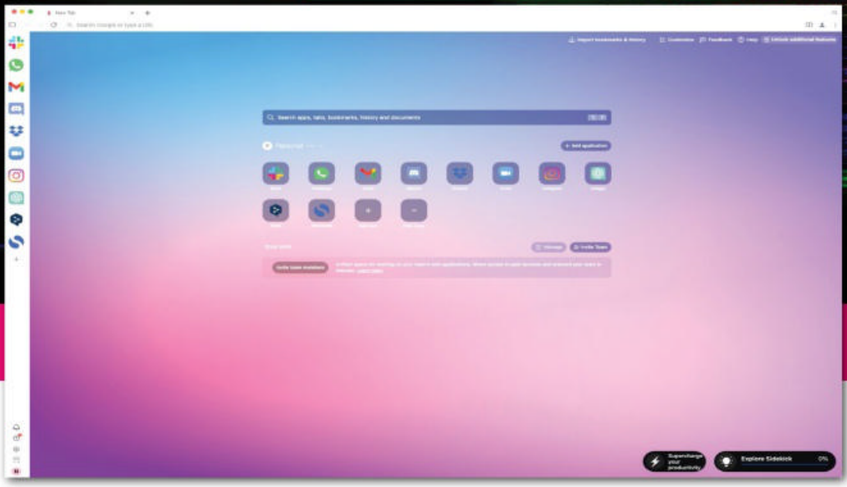

Sidekick aims to provide a distraction-free environment.

Sidekick is a browser with a purpose: it's designed to provide a distraction-free environment for work. Specifically, it's designed for people with ADHD, such as the browser's inventor, Dmitry Pushkarev, but that doesn't mean it should be pigeonholed as a browser for people with special needs.

"Whatever works for people with ADHD would also work for pretty much anyone else," said Pushkarev, who has a PhD in theoretical physics from Stanford and who co-founded Moleculo, a genome-decoding tech that was eventually acquired by Amazon.

Sidekick aims to keep users focused on the job in several ways. First, it attempts to save you from having to flick between different apps on the desktop by creating a sidebar of web apps that runs down the left of the browser screen. In the paid-for Pro version of the browser, those apps can be used in split-screen mode, allowing you to copy information from a Slack conversation and paste it into a Google Calendar appointment, say, without leaving the browser window.

At the bottom of the browser screen lurks a bell icon that allows you to switch off notifications for a set period of time or indefinitely, helping you to avoid becoming distracted by that new message indicator in your Gmail app. Likewise, ads, trackers, and other pop-up notifications are blocked by default to keep the browser clean and nimble. "It actually blocks all requests that are not important to render the page," said Pushkarev. "For example, if you open CNN, you will see that most of the requests are being blocked because those requests to trackers and ad networks are basically sending the data to third parties. We block these because you actually don't need to run them, and as a result we use a lot less memory and load the page faster.

"It doesn't seem like that big of a deal, an extra 100 milliseconds here and there," Pushkarev added. "But for people with ADHD, it's less chance to get distracted and makes you feel a lot less tired."

Tab management is also different to conventional browsers, something that's a common theme among all the power browsers in this feature. Sidekick groups tabs into "sessions"—so you might have a session for one work project and a separate session for another, or you might have a session of tabs that contain BBC Sport, Amazon, and other things that you typically browse during your lunch break. Again, the aim is to avoid losing focus.

I told Pushkarev about my little work tactic of keeping a couple of personal tabs open in the background, and then rewarding myself for five minutes every hour by reading soccer match reports or watching a YouTube video. He laughed, and then explained why that wouldn't work for people like him. "For people with ADHD, those five minutes can easily become three hours, and maybe even more, and then you suddenly find it's two in the morning and you're looking for some archaeological data, or you're on the Wikipedia article for Sudan. You're like, 'How did I get here?' and then you remember you actually started by trying to figure out some dashboards for the company. Not having those distractions is very, very useful, especially if you are easily distracted.

"For many people, keeping focused for even five minutes is a challenge," he added. "I know these people first-hand, and for them, life is virtually impossible without medication."

Earlier this year, the company added another feature to help people whose attention wanders. "We released a task

© SIDEKICK

SIDEKICK FACTS

Supported OSes: Windows, Mac, Linux
Browser engine: Chromium
Support for browser extensions: Yes (Chromium)
Founded: 2021

tracker with a built-in Pomodoro timer," he said, referring to the technique where you break down tasks into intervals of 25 minutes, followed by a short break. "Now you can say, 'Right now, I'm going to focus on this interview. I'm sticking with my screen saying [because] that's exactly what should be done at this point.'"

Sidekick offers a free version of the browser, but that's limited to five apps being pinned to the side of the screen, and has other limitations. The Pro version costs $10 per month, and there's a Team version that can be used within companies with pre-installed apps, built-in VPNs, and other collaboration features.

I put it to Pushkarev that convincing people to pay for a browser after they've been free for nearly 30 years will be a challenge. "It's probably your most important application," he replied. "This is the tool where you spend six, eight hours a day, and even small changes to this tool make a huge difference. We did a survey, and 80 percent of our users reported that they see substantial improvements to productivity when they use Sidekick."

If that proves to be the case, paying for a browser might not seem like such an extravagant idea after all.

© GETTY IMAGES

SESSIONS

Tabs are grouped into sessions, which can be renamed. Only those tabs open in the current session are visible in the tabs bar at the top of the screen, with the aim of stopping you checking Twitter updates. You can switch to different sessions from the sidebar or with a keyboard shortcut (Alt+W on Windows). Note the timer showing the current task coming to an end in the top-right of the browser window.

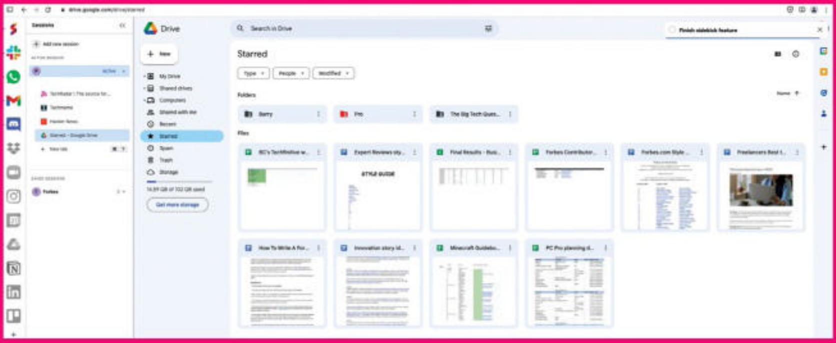

APPS SIDEBAR

Web apps are contained in a sidebar down the left-hand side of the browser window. Here, you'll still get notifications showing the number of unread Gmail or Slack messages, for example, but these can be easily turned off to avoid unwanted distractions. Apps don't get their own tab, so they take up less space at the top of the screen and avoid clutter.

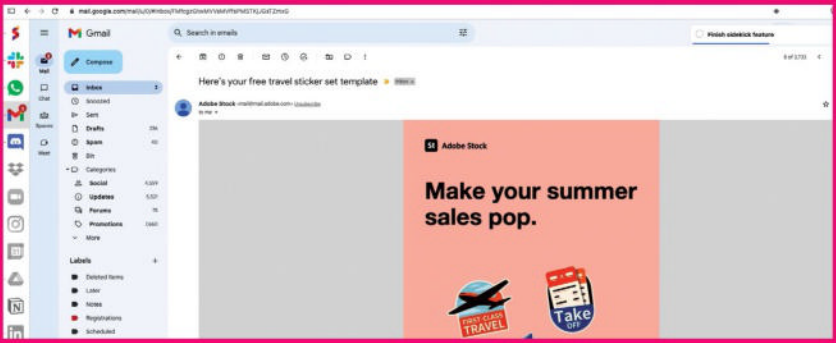

Here's your free travel sticker set template

Adobe Stock

Make your summer sales pop.

UNIVERSAL SEARCH

The browser's search menu, found by opening a new tab, can search across open tabs, apps, documents, and browser history to help you find things more quickly. As you can see from the screenshot, results appear instantly as you begin typing a search query, and it's easy to see which application the results are being sourced from.

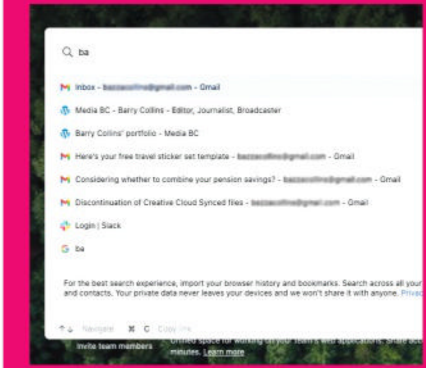

ba

Inbox - [redacted] - Gmail
Media BC - Barry Collins - Editor, Journalist, Broadcaster
Barry Collins' portfolio - Media BC
Here's your free travel sticker set template - [redacted] - Gmail
Considering whether to combine your pension savings? - [redacted] - Gmail
Discontinuation of Creative Cloud Synced files - [redacted] - Gmail
Login | Slack
ba

For the best search experience, import your browser history and bookmarks. Search across all your and contacts. Your private data never leaves your devices and we won't share it with anyone. Privac

DISTRACTION BLOCKER

The Distraction Blocker can be activated manually or according to a schedule. So you might have it switched on between 9am and 5pm on weekdays, when you're meant to be focused on work. It can be set to automatically redirect you away from time-sucking websites. Those distractions and redirects can be customized, in case you have a particular weakness for the websites of soccer clubs during work hours, say.

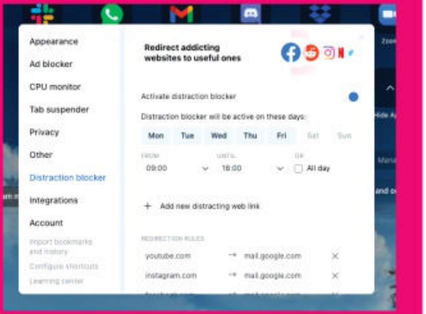

Redirect addicting websites to useful ones

Activate distraction blocker ●

Distraction blocker will be active on these days:

Mon Tue Wed Thu Fri Sat Sun

FROM 09:00 UNTIL 18:00 OR ☐ All day

+ Add new distracting web link

REDIRECTION RULES

youtube.com → mail.google.com ✕
instagram.com → mail.google.com ✕

VIVALDI

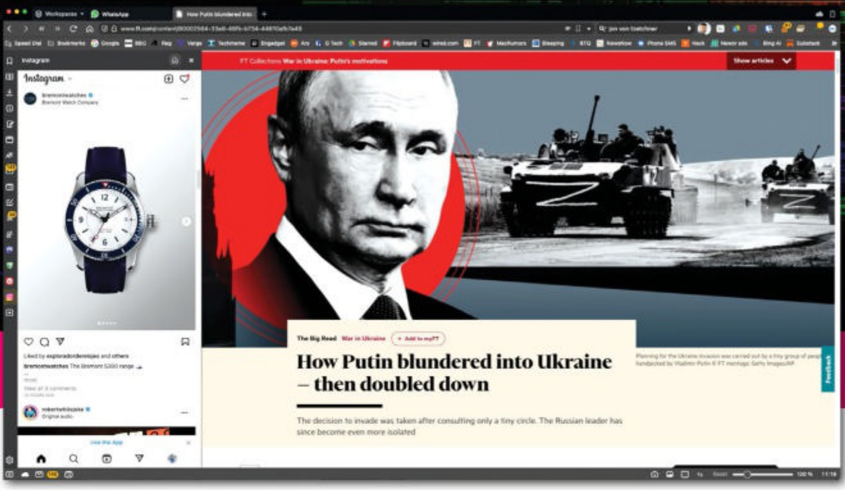

How Putin blundered into Ukraine — then doubled down

The decision to invade was taken after consulting only a tiny circle. The Russian leader has since become even more isolated

Feature-packed Vivaldi relies heavily on user feedback.

With a full-blown email client, calendar with offline access, RSS reader, task manager, notes app, and lots more available through extensions, it's tempting to ask whether Vivaldi is more of an operating system than a web browser.

The company isn't blind to the complexity of its offering. When you first install Vivaldi, you're given three layout options: Essentials, Classic, and Fully Loaded. Essentials hides pretty much everything bar the basic browsing tools; Classic adds a few more advanced features, such as Web Panels (see opposite), and the browser's fast-forward and rewind buttons; Fully Loaded gives you the lot.

But if you're looking for an apology from Vivaldi CEO, Jon von Tetzchner, for packing the browser with features, you'd best bring a book; you'll be waiting a long time. He launched Vivaldi in 2015 with the clear idea of making a highly customizable, privacy-first browser that stood apart from the feature-stripped mainstream browsers of the time. Getting on for ten years later, he's created a browser that's different from anything else on the market, even if he now believes that others are "borrowing" Vivaldi's ideas.

"If you look at a lot of the features that are being added by our competitors, and even some of the new ones, a number of them are basically features we have that they've taken and maybe made some changes to," von Tetzchner told us. "I'm not saying we innovated everything and everything is basically ours, but it is interesting. Some are adding [features] like tab tiling, and we've been doing that for quite some time. Suddenly, a number of the browsers have been adding panels on the right side or they're providing vertical tabs. We were the only ones that were doing those things."

How does Vivaldi decide which features to add to its browser? "It's a process where we work with the user base quite closely," said von Tetzchner. "Some of the ideas are just people asking for features—that happens. Typically that would be features that you would find in old versions of Opera," he added, Opera being the browser he co-founded in the early 1990s, before falling out with the browser's new owners in 2011.

"Then there are people having ideas internally," von Tetzchner added. "We have a concept of 'brains think better than the brain'. So anyone could come up with an idea, and sometimes it's a developer that just implements something and they say, 'Hey, I made this prototype, what do you think?' If there's agreement that it's a good idea, then we go and implement it further, so it's less structured. I think it's very difficult to plan everything."

Vivaldi relies heavily on feedback from user groups when implementing new features. While von Tetzchner concedes that there are dangers in relying too heavily on the feedback of a group of super-engaged, technically literate beta testers, he said there are big advantages from listening carefully to that group, too. "We are always getting feedback from those natural users," he said. "Yes, they are more advanced than the rest, but at the same time, we are building a browser that probably some of the more advanced users will really appreciate, because we have a lot of features that you don't find anywhere else. So, from that perspective, this is a very relevant group for us to test."

However, there will inevitably come a time with such a feature-packed browser that features have to be removed to keep the code base lean and manageable. Vivaldi makes a big play of not tracking its users' behavior, so how does the company decide which features to eliminate when it doesn't have

© VIVALDI

VIVALDI FACTS

Supported OSes: Windows, Mac, Linux, Android, iOS (in beta)
Browser engine: Chromium
Support for browser extensions: Yes (Chromium)
Founded: 2015

a barrage of analytics to be sure that a certain feature isn't widely used? Von Tetzchner admits that it's difficult.

"Typically, when we try to remove [a feature], we actually find out it's the favorite feature for a certain group of people," he said. "Which is why we then end up removing it as a default feature, but keeping it around as an option, which doesn't need to be that expensive a thing to do. But obviously, is there a limit? Yes. Sometimes we find that we can actually remove features if we're providing a better way to do things, which everyone agrees is better."

It's not only the feature set of Vivaldi that's constantly expanding; it's the number of devices it supports, too. Aside from the desktop browser for Windows, Mac, and Linux, it launched an Android version in 2020, and this is also widely used to provide browsing capabilities for in-car systems. At the time of writing, Vivaldi was gearing up to launch an iOS version as well, having long resisted Apple's demands to use its WebKit framework. "We don't get the benefit of having the same codebase," said von Tetzchner, "but we're still trying to make sure that the functionality that you're used to having is there and it's recognizable." ⏻

© GETTY IMAGES

EMAIL CLIENT

Vivaldi's email app isn't merely a shortcut to webmail, which you might find in other browsers. It's a full-blown email client for webmail accounts, as well as IMAP/POP3. It offers full offline access, filters, search, and support for multiple accounts. Notifications of unread messages appear in the sidebar, and you can alter the frequency for which it checks for new messages, among many other settings.

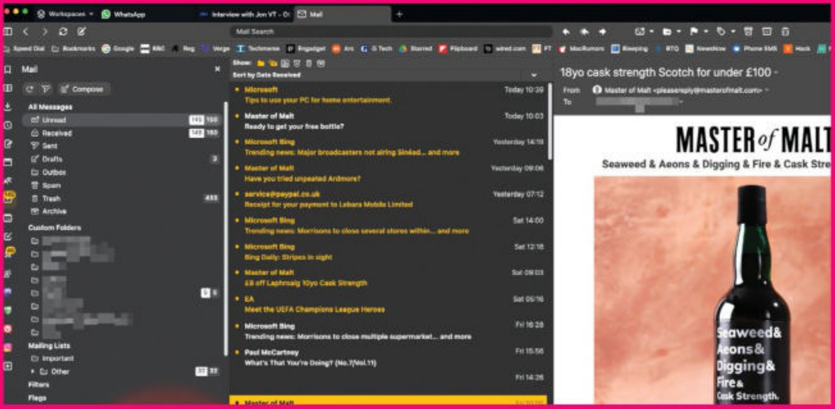

Microsoft — Today 10:39
Tips to use your PC for home entertainment.

Master of Malt — Today 10:03
Ready to get your free bottle?

Microsoft Bing — Yesterday 14:19
Trending news: Major broadcasters not airing Sinéad... and more

Master of Malt — Yesterday 09:06
Have you tried unpeated Ardmore?

service@paypal.co.uk — Yesterday 07:12
Receipt for your payment to Lebara Mobile Limited

Microsoft Bing — Sat 14:00
Trending news: Morrisons to close several stores within... and more

Microsoft Bing — Sat 12:18
Bing Daily: Stripes in sight

Master of Malt — Sat 09:03
£8 off Laphroaig 10yo Cask Strength

EA — Sat 05:16
Meet the UEFA Champions League Heroes

Microsoft Bing — Fri 16:28
Trending news: Morrisons to close multiple supermarket... and more

Paul McCartney — Fri 15:56
What's That You're Doing? (No.7/Vol.11)

18yo cask strength Scotch for under £100
From Master of Malt <pleasereply@masterofmalt.com>

RSS READER

RSS feeds are still a convenient way of keeping on top of updates from websites, newsletters, and suchlike. Vivaldi's RSS reader is a little awkwardly integrated with the email client, although if you fiddle with the settings, you can keep them separate. You can choose how often Vivaldi checks for new updates when you subscribe to a feed, with options ranging from every five minutes to weekly.

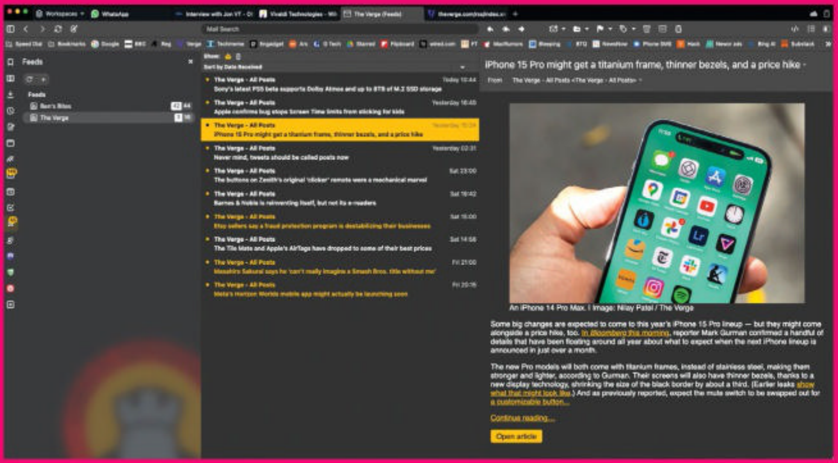

The Verge - All Posts — Today 10:44
Sony's latest PS5 beta supports Dolby Atmos and up to 8TB of M.2 SSD storage

The Verge - All Posts — Yesterday 16:40
Apple confirms bug stops Screen Time limits from sticking for kids

The Verge - All Posts
iPhone 15 Pro might get a titanium frame, thinner bezels, and a price hike

The Verge - All Posts — Yesterday 02:31
Never mind, tweets should be called posts now

The Verge - All Posts — Sat 23:00
The buttons on Zenith's original 'clicker' remote were a mechanical marvel

The Verge - All Posts — Sat 19:42
Barnes & Noble is reinventing itself, but not its e-readers

The Verge - All Posts — Sat 15:00
Etsy sellers say a fraud protection program is destabilizing their businesses

The Verge - All Posts — Sat 14:58
The Tile Mate and Apple's AirTags have dropped to some of their best prices

The Verge - All Posts — Fri 21:00
Masahiro Sakurai says he 'can't really imagine a Smash Bros. title without me'

The Verge - All Posts — Fri 20:16
Meta's Horizon Worlds mobile app might actually be launching soon

iPhone 15 Pro might get a titanium frame, thinner bezels, and a price hike
From The Verge - All Posts <The Verge - All Posts>

An iPhone 14 Pro Max. | Image: Nilay Patel / The Verge

Some big changes are expected to come to this year's iPhone 15 Pro lineup — but they might come alongside a price hike, too. In Bloomberg this morning, reporter Mark Gurman confirmed a handful of details that have been floating around all year about what to expect when the next iPhone lineup is announced in just over a month.

The new Pro models will both come with titanium frames, instead of stainless steel, making them stronger and lighter, according to Gurman. Their screens will also have thinner bezels, thanks to a new display technology, shrinking the size of the black border by about a third. (Earlier leaks show what that might look like.) And as previously reported, expect the mute switch to be swapped out for a customizable button...

WORKSPACES

Workspaces basically bring the concept of virtual desktops to browsers, where each workspace gets its own set of tabs (think Work, Shopping, Holiday), allowing you to keep groups of tabs isolated from another. Workspaces can be easily renamed and assigned a custom icon.

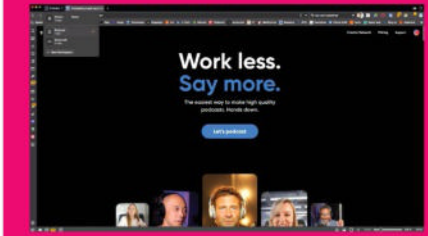

WEB PANELS

Web Panels create a pane at the side of the screen that you can use to, say, keep an eye on soccer results, stock prices, or social media while you get on with stuff in the main window. They're normally treated like mobile browsers, so content is reformatted to fit the pane.

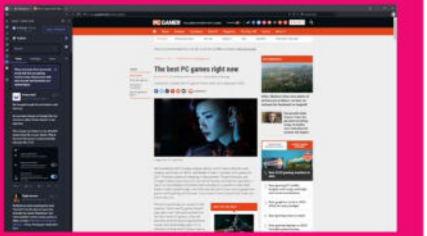

BUILD YOUR OWN MUSIC SERVER FOR FREE

Nik Rawlinson explains how to break free from Spotify's monthly bills by running your own free music server that works with Alexa

© MILAN MARKOVIC/E+/GETTY IMAGES

ONE OF THE BEST USES for a smart speaker is as an on-demand jukebox. Amazon's Echo devices integrate directly with both Spotify and Amazon's own online music service, each of which offers a library of more than 100 million tracks, so you should be able to find almost everything you want.

There are downsides, however. One is that you need to pay a subscription fee to get the best from these music services. On both Spotify and Amazon Music, free users can only request playlists, rather than specific songs; Amazon Prime subscribers can stream individual tracks from Amazon Music, but the service limits how often you can skip, and insists on shuffling albums and playlists. For full control you need to pay $10.99 a month for a full Spotify account, or $10.99 a month for Amazon Music Unlimited.

Another catch is that, as we said, you can find almost everything you want on these services. But what about those rare bootlegs and B-sides that you've been painstakingly collecting? Or recordings of your friends' bands, and your own mixes? Curating your own music library is unfashionable these days, but it's the best way to keep full control of your music, with all the tracks you want to listen to and no recurring payments.

The good news is that it's perfectly possible to get Alexa to stream your own music: you simply need to use the right software to make your tracks available over the network. A great free solution is Plex Media Server (and it also happens to handle video streaming as well). Here's how to set up Plex in Windows, and—our preferred solution—how to run it on an inexpensive, low-power Raspberry Pi computer.

SETTING UP PLEX MEDIA SERVER ON WINDOWS

The first thing you need is a free Plex account. To create one, point your browser at plex.tv and click "Sign up". Log in with Google, Facebook, or Apple, or create a Plex account with your email address and a password; then click "Create an Account". You can subscribe to Plex's premium services at a later date if you wish, but for our purposes, a free account is fine.

© PLEX

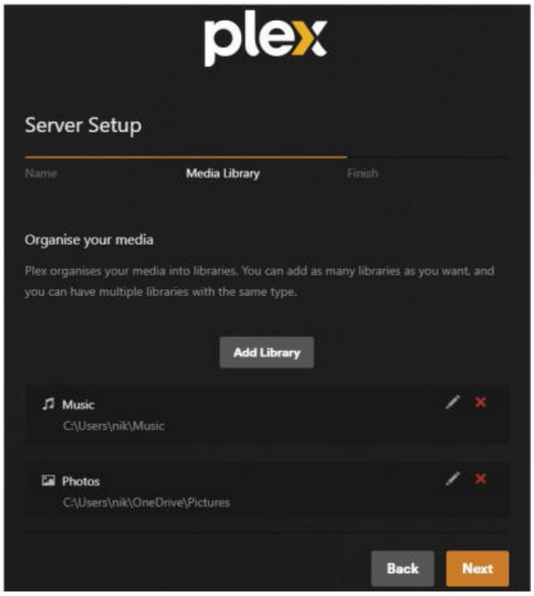

Server Setup

Organise your media

Plex organises your media into libraries. You can add as many libraries as you want, and you can have multiple libraries with the same type.

Music
C:\Users\nik\Music

Photos
C:\Users\nik\OneDrive\Pictures

You can create multiple libraries for various media types.

As well as music services, Plex integrates with a range of video-streaming platforms, including Amazon Prime, Apple TV+, and Netflix. When you first sign in, it offers to add them to your player; you can connect them to your account by selecting them in the list, or by closing the streaming services dialog box.

Having created your account, you need to install the server software, which runs in the background and makes your tracks available to playback devices. Click "+ Your Media" in the sidebar, followed by "Get Plex Media Server" . You can also install the server application from the Microsoft Store at tinyurl.com/2wwhcby5.

Once the download completes, step through the installation process and start the server. The setup wizard then launches to walk you through the process of adding media to your library. This doesn't involve actually moving any files, just adding their details to Plex's database.

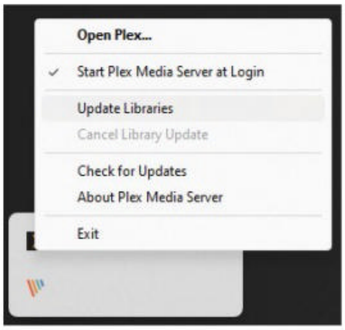

If your libraries aren't refreshing, force an update from the taskbar icon.

The setup wizard will also ask if you want to synchronize your watch state and recordings. If you do, any ratings you apply to songs, TV shows, or movies will be replicated across all devices logged in to the same account. It's up to you what you pick here.

The next screen explains how the system works; click "Got It", then give your library a name. By default, Plex uses the name of your computer, but if several people use the same PC, it might make sense to use your own name.

By default, Plex tries to make your media available not just over your home network, but across the internet, so you can listen on any device, anywhere in the world, as long as it's signed into your Plex account. Leave this option ticked; if you later find that external access isn't working, check that your firewall allows access to the main Plex port, which is TCP port 32400. Other ports worth checking include UDP ports 1900, 5353, 32410, 32412, 32413, and 32414, and TCP ports 8324 and 32469. For further details, Plex's help page on the subject is at tinyurl.com/yp47nhwt.

Now, click Next, and decide which libraries you want to make accessible through your server. Music and Photos may be selected by default, and already pointing to the music and photos folders on your PC. You can get rid of any libraries you don't want by clicking "X" to remove them, or edit their media locations by clicking the pencil icon.

If you want to share other types of media, click Add Library and navigate to the folder where the items are located. Don't worry too much about the contents of your media folders at this point—you can add as much content as you like later on. Once you're happy with your Library configuration, click Next, followed by Done.

Plex Media Server will now start cataloguing your media. To browse the music it's discovered, click More in the sidebar, followed by the Music link at the top of the following sidebar. (The Music link further down takes you to the Tidal streaming service, rather than your own music.)

If your library shows as empty, it may simply be that Plex hasn't scanned it yet. To force it to refresh

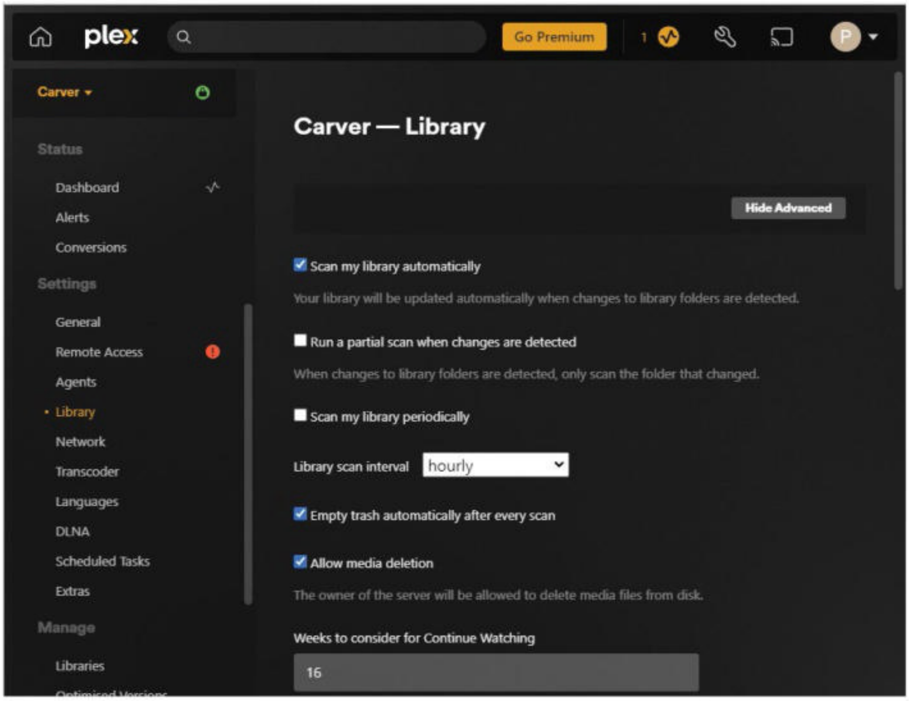

Carver — Library

☑ Scan my library automatically

Your library will be updated automatically when changes to library folders are detected.

☐ Run a partial scan when changes are detected

When changes to library folders are detected, only scan the folder that changed.

☐ Scan my library periodically

Library scan interval: hourly

☑ Empty trash automatically after every scan

☑ Allow media deletion

The owner of the server will be allowed to delete media files from disk.

Weeks to consider for Continue Watching: 16

Ensure Plex is set up to refresh your libraries automatically.

its database, click the upward arrow in the Windows system tray to reveal your hidden icons; now right-click the Plex icon and select Update Libraries from the context menu. Plex will refresh its listing from each of your libraries and add what it finds to the Plex Media Server, with albums represented by their cover art.

Another possibility is that Plex hasn't been able to recognize the tracks in your library. The software does its best to identify music files from their paths and filenames, and can even read embedded ID3, but you'll get the best results if you organize your tracks according to the preferred Plex arrangement. This uses a separate folder for each artist, with a subfolder for each album, EP, or single, containing numbered tracks in MP3, AAC, FLAC, M4A, or WAV format. Your library will probably contain something like the following:

Music\Toto\Toto IV\01—Rosanna.mp3
Music\Toto\Toto IV\02—Make Believe.mp3
Music\Toto\Toto IV\03—I Won't Hold You Back.mp3

and so forth. You can find more information about Plex's file-naming conventions at tinyurl.com/3kjscwsv.

It's also worth making sure that the system is set to scan your library locations for updates, so that any new files you add will become available in Plex. To do this, click the Activity icon at the top of the screen (shown as a pulse inside an orange circle), then scroll down to the section that shows your server name and click Library. Make sure the box beside "Scan my library automatically" is ticked. To make sure nothing's been missed, you can also tick the box beside "Scan my library periodically" and select an interval from the dropdown menu on the line below.

You can find and play music files from any web browser by connecting to the computer running the server software on port 32400. For a computer inside your home network, the address might be something like http://10.0.0.44:32400.

You can also access the web interface over the internet; you simply need to ensure that your router is forwarding TCP port 32400 to the system running your Plex server. Either browse to the external IP address of your router, or visit plex.tv, log in, and select "Open Plex" from the dropdown. You can also connect via the Plex mobile app for Android and iOS.

Playing songs with Alexa is possible, but first you'll need to link your Plex and Alexa accounts (see "Connect your Plex account to Alexa"). There's no official way to stream from Plex to a Google Home device using voice commands, but you can use the Plex mobile app to stream music to any Chromecast-compatible receiver.

SETTING UP PLEX MEDIA SERVER ON A RASPBERRY PI

It may be convenient to run Plex on your desktop, but your songs will only be available while the computer's powered on and awake. If you want an always-on server, that's ideal for a Raspberry Pi. You don't even need a recent model if you only want to serve up audio, although high-definition video is likely to overtax even the most powerful board. You may, however, want to use a USB hard drive to store your media files instead of squeezing them onto an SD card.

The easiest way to get started with Plex on the Pi is to install the DietPi OS build (dietpi.com), rather than the default Raspberry Pi OS. This optimized version of Debian is preconfigured to run a wide range of services, including Nextcloud, Pi-hole, Kodi, WordPress, and Plex.

This means you can get set up in just a few simple steps. Start by connecting your external drive to the Pi, then booting the DietPi image. Next, go back to your Windows machine, open a Command Prompt, and connect to your Raspberry Pi via ssh, using the username "dietpi" and the board's IP address. In our case, the command is:

```
ssh dietpi@10.0.0.50
```

Rather than filling up your SD card, attach an external hard drive.

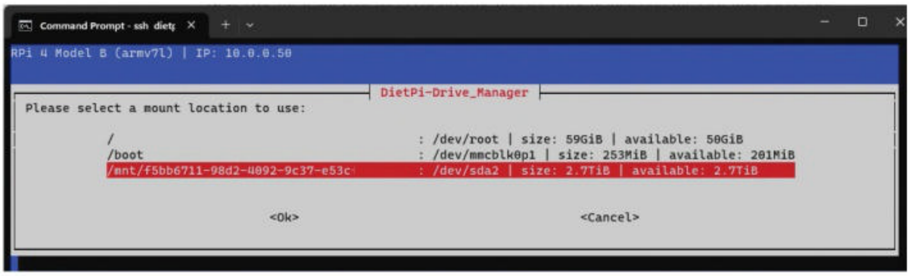

```
RPi 4 Model B (armv7l) | IP: 10.0.0.50

                         DietPi-Drive_Manager
Please select a mount location to use:

      /                                  : /dev/root | size: 59GiB | available: 50GiB
      /boot                              : /dev/mmcblk0p1 | size: 253MiB | available: 201MiB
      /mnt/f5bb6711-98d2-4092-9c37-e53c  : /dev/sda2 | size: 2.7TiB | available: 2.7TiB
```

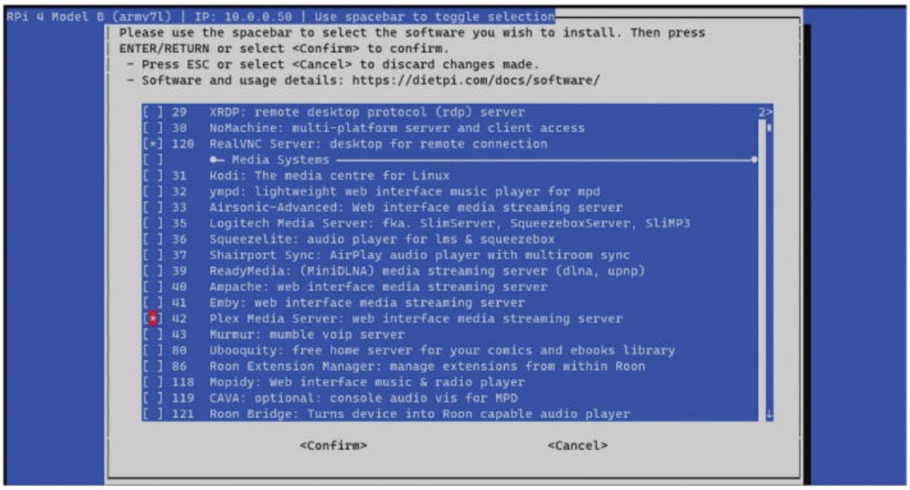

```
RPi 4 Model B (armv7l) | IP: 10.0.0.50 | Use spacebar to toggle selection
 Please use the spacebar to select the software you wish to install. Then press
 ENTER/RETURN or select <Confirm> to confirm.
  - Press ESC or select <Cancel> to discard changes made.
  - Software and usage details: https://dietpi.com/docs/software/

   [ ] 29   XRDP: remote desktop protocol (rdp) server
   [ ] 30   NoMachine: multi-platform server and client access
   [*] 120  RealVNC Server: desktop for remote connection
   [ ]      ●─ Media Systems ───────────────────────────────────────●
   [ ] 31   Kodi: The media centre for Linux
   [ ] 32   ympd: lightweight web interface music player for mpd
   [ ] 33   Airsonic-Advanced: Web interface media streaming server
   [ ] 35   Logitech Media Server: fka. SlimServer, SqueezeboxServer, SliMP3
   [ ] 36   Squeezelite: audio player for lms & squeezebox
   [ ] 37   Shairport Sync: AirPlay audio player with multiroom sync
   [ ] 39   ReadyMedia: (MiniDLNA) media streaming server (dlna, upnp)
   [ ] 40   Ampache: web interface media streaming server
   [ ] 41   Emby: web interface media streaming server
   [*] 42   Plex Media Server: web interface media streaming server
   [ ] 43   Murmur: mumble voip server
   [ ] 80   Ubooquity: free home server for your comics and ebooks library
   [ ] 86   Roon Extension Manager: manage extensions from within Roon
   [ ] 118  Mopidy: Web interface music & radio player
   [ ] 119  CAVA: optional: console audio vis for MPD
   [ ] 121  Roon Bridge: Turns device into Roon capable audio player

              <Confirm>                          <Cancel>
```

Plex Media Server is included in the DietPi application bundle.

Once you've logged in, you'll see a Linux command prompt. Enter:

```
sudo dietpi-software
```

First, configure the Pi to use the USB drive for data storage, rather than the internal SD card. Key down to "User Data Location" and press Return, select "List", and press Return again. Select your attached drive and press Return. This records the change and returns you to the previous screen.

Next, key up to "Browse Software" and press Return to open the full list of applications that come with DietPi. Scroll until you come to item "42: Plex Media Server". Press Space to select it.

While we're here, let's install DietPi's file-management tool, as this will make it easier to organize your media library. Scroll until you get to "198: File Browser", and select it by pressing Space. Press Tab to highlight "Confirm", and press Return to go back to the previous screen. To confirm the installations, key down to "Install" and press Tab (which selects "OK"), then press Return.

Once this is done, you should be able to log into Plex via a browser in just the same way as on a PC. Open a browser on any other computer on your network, and enter the IP address of the board, with ":32400" appended to specify the Plex port. When you get there, sign in using your existing Plex credentials.

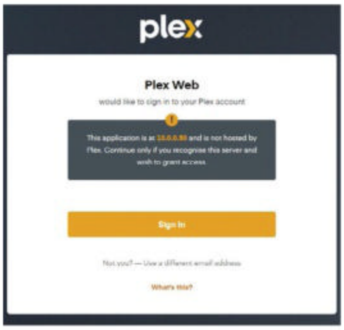

Plex Web

would like to sign in to your Plex account

This application is at 10.0.0.50 and is not hosted by Plex. Continue only if you recognise this server and wish to grant access.

The media server may warn that it's not being hosted by Plex.

© PLEX

You'll now be walked through the same setup process as outlined above. There's one key difference, though. When creating a media library, you'll be presented with a "Browse for Media Folder" button; don't click this. Instead, click in the address box and enter:

```
/mnt/dietpi_userdata/Music/
```

This is the address of the default Music folder that's created by the File Browser as part of its installation process. Note the capital M, as Linux-type operating systems are case-sensitive.

You can now click through the installer until you reach the Plex interface. Click the activity icon at the top and select your server; click "Library" in the sidebar, and make sure "Scan my library automatically" is ticked. Scroll to the bottom and click Save Changes.

All that remains is to copy your files to the USB. Open a browser on your Windows machine and navigate to your Pi's IP address, followed by port 8084—for us, that's 10.0.0.50:8084. Log in using the username dietpi.

Now, click into the dietpi_userdata folder, and from there into Music. Click the up-pointing arrow at the top of the window, followed by Folder, to upload a complete folder. Select the folder you want to upload to the server, then click Upload to confirm the transfer.

If Plex doesn't immediately catalog your tracks, return to the Plex Media Server interface, hover over your music library in the sidebar, and click the three vertical dots to the right of its name. Click "Scan Library Files" on the fly-out menu. When Plex has updated its database, you should be able to play your tracks from the web interface, mobile app—or, once you've linked your accounts, through Alexa.

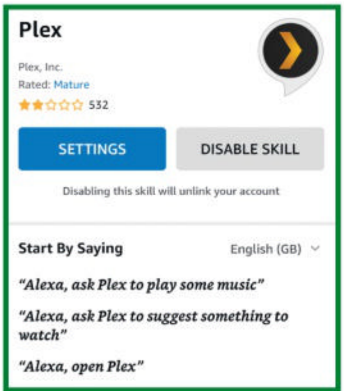

Plex

Plex, Inc.

Rated: Mature

★★☆☆☆ 532

Disabling this skill will unlink your account

Start By Saying

English (GB)

"Alexa, ask Plex to play some music"

"Alexa, ask Plex to suggest something to watch"

"Alexa, open Plex"

CONNECT YOUR PLEX ACCOUNT TO ALEXA

Once you've got your Plex server running, you need to connect it to your Alexa account to make its content available. Open the Alexa app on your phone and tap More, followed by "Skills & Games". Use the box at the top of the screen to search for "Plex". Tap the result, then tap the "Enable to Use" button.

The app will switch to your browser, load the Plex website, and prompt you to log in. Once you've done this, you'll be asked to authorize Amazon Alexa to access your Plex account; tap "Authorize", followed by "Close".

You can now close the Alexa app, and start streaming your songs, albums and playlists by saying "Alexa, ask Plex to play..." Your requests are passed straight through to Plex, so any new files or playlist updates will be instantly available as soon as they're indexed.

The Electronic Frontier Foundation is the leading nonprofit organization defending civil liberties in the digital world. Founded in 1990, EFF champions user privacy, free expression, and innovation through impact litigation, policy analysis, grassroots activism, and technology development. We work to ensure that rights and freedoms are enhanced and protected as our use of technology grows.

EFF.ORG

ELECTRONIC FRONTIER FOUNDATION

Protecting Rights and Promoting Freedom on the Electronic Frontier

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

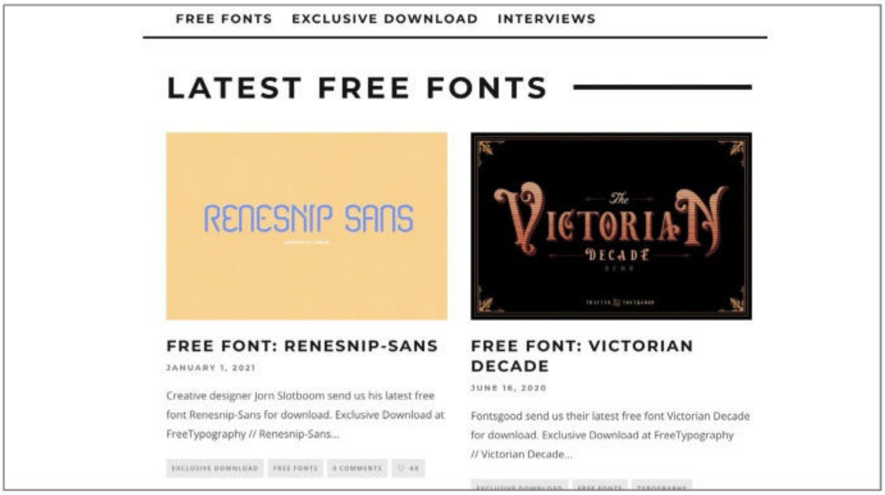

FREE FONTS EXCLUSIVE DOWNLOAD INTERVIEWS

LATEST FREE FONTS

FREE FONT: RENESNIP-SANS

JANUARY 1, 2021

Creative designer Jorn Slotboom send us his latest free font Renesnip-Sans for download. Exclusive Download at FreeTypography // Renesnip-Sans...

EXCLUSIVE DOWNLOAD FREE FONTS 0 COMMENTS ♡ 40

FREE FONT: VICTORIAN DECADE

JUNE 16, 2020

Fontsgood send us their latest free font Victorian Decade for download. Exclusive Download at FreeTypography // Victorian Decade...

EXCLUSIVE DOWNLOAD FREE FONTS TYPOGRAPHY

TYPE, TYPE, AND TYPE

When creating something, people tend to envisage the overall look and feel of a specific design and the typography is left till the end. Well, it certainly shouldn't be, and a great place to head to for free fonts and inspiration is www.freetypography.com. There are pages upon pages of examples with mockups included to see what your text will look like. It's a great place to start your design, instead of making it an afterthought.

MAKE – USE – CREATE

64
Autopsy: Apple iPhone 15 Pro Max

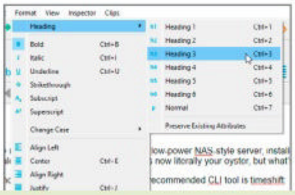

66
Create and export your own eBooks

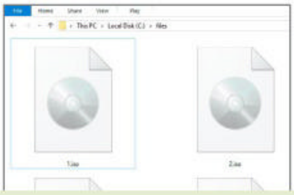

70
Make secure backups using an ISO file

© IFIXIT, SIGIL-EBOOK.COM, WINCDEMU

SAM LEWIS
STAFF WRITER

IT'S HAPPENING...

Last time, I mentioned DLSS 3.5, an AI powerhouse from Nvidia that was bound to take over my life and keep me glued to my PC. Well, here we are—it didn't take long for it to wow me. I've become the office gremlin that I once feared. Okay, it's not that bad—I don't have piles of food wrappers on my desk, and I have been outside, so it's fine, alright?

Not only do we have access to Nvidia DLSS 3.5, but one of my favorite titles to take advantage of this AI feature, *Cyberpunk 2077*, has an expansion, which I'm excited is already using DLSS 3.5. So how do I balance this with life? I'm sure most of you know the struggle of being hooked on a new game or DLC, and with AI helping out as much as it does by making an already pretty game even more beautiful, I can't help but be immersed in the experience.

At the time of writing, the expansion pack is due to be released in four days. Named *Phantom Liberty*, it's a major expansion that adds more stories, quests, collectibles, and abilities. It will also star one of my favorite actors, Idris Elba, as a playable character. It's essentially a spy-thriller-esque story but all this fun doesn't come for free. With a $30 price tag, you'll have to apologize to your wallet and admit defeat just like I have.

↘ submit your How To project idea to: editor@maximumpc.com

IFIXIT presents:

AUTOPSY

THIS MONTH WE DISSECT...

Apple iPhone

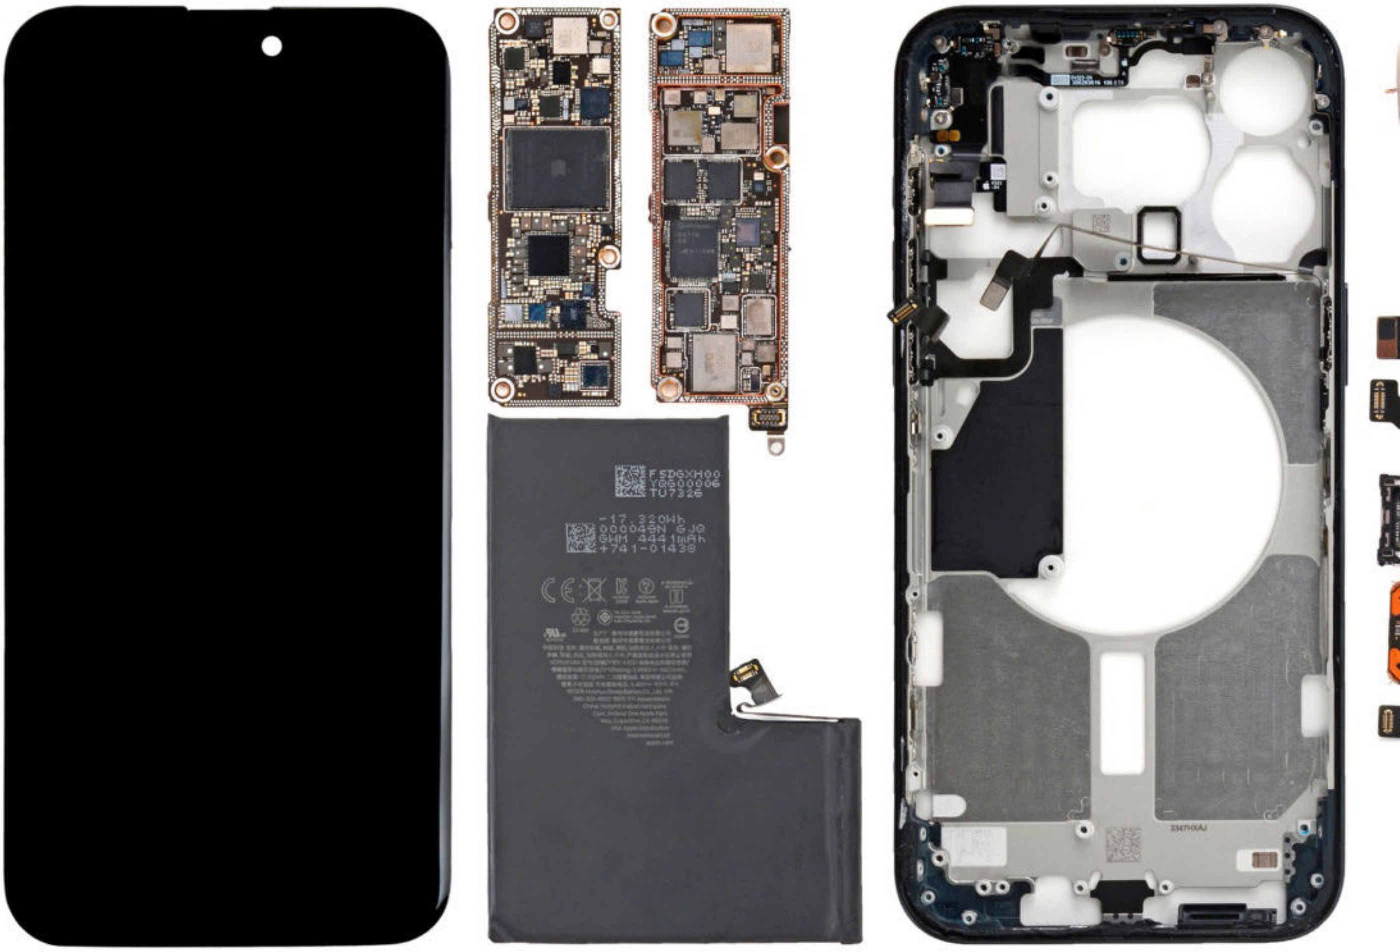

About iFixit

iFixit is a global community of tinkerers dedicated to helping people fix things through free online repair manuals and teardowns. iFixit believes that everyone has the right to maintain and repair their own products. To learn more, visit www.ifixit.com

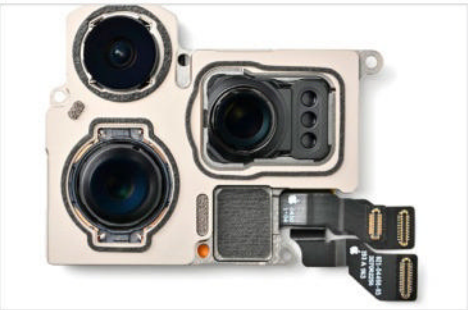

Instead of lens elements controlled by electromagnets, Apple designed a "periscope" simulating a 120mm focal length.

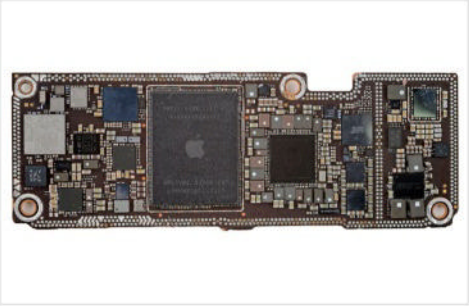

The Apple A17 Pro hexa-core applications processor with six-core GPU layered under what is likely SK Hynix DDR5 memory.

All of the internals are hiding behind the screen instead of behind the back glass. However, that back glass is now removable, too, just like the 14.

15 Pro Max

You can use an iPhone 15 Pro to charge AirPods, Apple Watch, or another device that supports USB Power Delivery at up to 4.5 watts.

The battery in the 15 Pro Max is 4422 mAh, a 2.3 percent increase over the 14 Pro Max's 4323 mAh pack.

BACKGROUND

iPhone 15 Pro Max is here, and it boasts a new charge port, titanium frame, periscope, and new pro-repair design. Let's dive inside!

MAJOR TECH SPECS

- A17 Pro SoC with 6-core CPU, 6-core GPU and 16-core neural engine
- 6.7-inch Super retina XDR display, 2,796 x 1,290 pixels, 2,000 nits
- 48MP, 24mm, *f*/1.78 aperture main camera; 12MP, 13mm, *f*/2.2 aperture ultra-wide camera; 12MP, 120mm, *f*/2.8 aperture, 5x telephoto camera
- 8GB of RAM and 256GB of storage (configurable up to 1TB)
- 5G (sub-6 GHz) with 4x4 MIMO7, Gigabit LTE with 4x4 MIMO and LAA7, Wi-Fi 6E with 2x2 MIMO8, Bluetooth 5.3
- MagSafe 15W wireless charging
- IP68 water-resistance rating

KEY FINDINGS

- The 15 Pro and Pro Max now open from both directions, but in an unexpectedly opposite manner: all of the internals are behind the screen instead of the back glass. But that back glass is removable, too. This inverted chassis arrangement makes critical repairs riskier than on the 14, because you're removing the fragile display rather than an inert sheet of glass.

- To get to the battery, you have to heat and pry open the screen, which is riskier. If you accidentally tear a cable, it's better to tear the back cover cable than ruin the display.

- The battery is serviceable only by a trained technician. Fortunately, you can be one! Simply follow Apple's service manual or iFixit's repair guides. It's not all that hard—we've helped millions of people fix their electronics themselves.

- Once we separated the main boards, we can confirm that Apple is using Broadcom's high-end Snapdragon X70 modem. According to Qualcomm, it's powered by AI for beam management and antenna tuning. Alas, 8K isn't present.

- The big upgrade for the cameras is the "tetraprism" periscope lens, upping the optical zoom from 2x to 5x. The S23 Ultra technically beats this with 10x zoom, but the way Apple achieved it is interesting.

- Aside from the new periscope lens, the sensors on the 15 Pro Max main and wide cameras seem to be the same size as last year's 14 Pro Max—suggesting that any improvement in image quality has more to do with the new A17 SoC than the hardware.

- Unfortunately, swapping the front camera between two of our 15 Pro Max units results in complete failure. When we first encountered this problem, we assumed that it was a bug that would be quickly fixed. Our trust was clearly misplaced.

- Parts pairing in these models extends beyond mechanical compatibility, requiring authentication and pairing through Apple's System Configuration tool, limiting genuine replacements to Apple-blessed ones and impacting independent repair enterprises and the issue of e-waste.

- In addition to compatibility advantages, the new USB-C port can provide 4.5 watts of power to external devices. That's a 15x upgrade over previous Lightning phones, which can only output 0.3 watts.

- Repairability Score: 4 out of 10 (10 is easiest to repair). The 15 is an incremental upgrade. The rollout of the 14's open-from-both-sides architecture is welcome. Investments in camera hardware are always well spent, and the periscope particularly stands out. Unfortunately, software is the anchor around an otherwise exceptionally designed phone, but without the ability to swap components, repairability suffers dramatically. ⏻

Create and export your own eBooks

YOU'LL NEED THIS

SIGIL
(https://sigil-ebook.com)

EPUB EDITOR
and its companion app Page Edit to allow you to format your book's text using a WYSIWYG editor instead of XHTML code.

EVER FANCIED YOURSELF AS AN AUTHOR? These days, you don't need a publishing deal to distribute your magnum opus, and you don't even need to commit to the upfront cost of printing copies without any guaranteed sales in place. Thanks to the rise of eReaders like the Kobo and Kindle, you can publish digitally for zero cost and then look to promote your work online.

The hardest part of the job is taking your finished draft—typically in a Microsoft Word document—and converting it into a format that opens on an eReader flows properly on any screen and is easy to read. The solution lies in an open-source program called Sigil, which takes your text (and any images you may have) and transforms it into a standard ePub file, which can be read on a variety of mobile devices, from Kobo eReaders to tablets and even phones.

At first glance, Sigil can appear impenetrable, but this tutorial will help you get to grips with the basics, allowing you to create and export your book in a format you can also easily convert to the AZW3 format used by Kindles for maximum exposure. –**NICK PEERS**

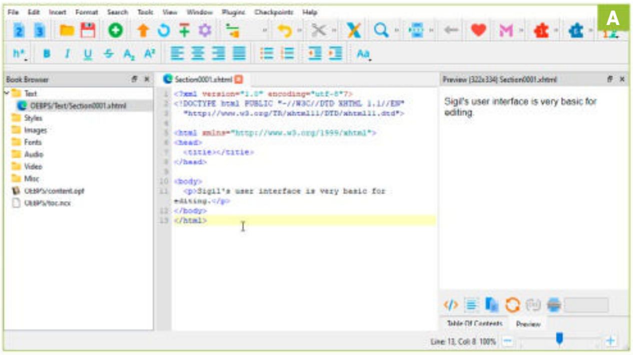
A
File Edit Insert Format Search Tools View Window Plugins Checkpoints Help
Book Browser
Text
OEBPS/Text/Section0001.xhtml
Styles
Images
Fonts
Audio
Video
Misc
OEBPS/content.opf
OEBPS/toc.ncx
Section0001.xhtml
Preview (322x334) Section0001.xhtml
Table Of Contents
Preview
Line: 13, Col: 8 100%

```
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.1//EN"
  "http://www.w3.org/TR/xhtml11/DTD/xhtml11.dtd">

<html xmlns="http://www.w3.org/1999/xhtml">
<head>
  <title></title>
</head>

<body>
  <p>Sigil's user interface is very basic for
editing.</p>
</body>
</html>
```

Sigil's user interface is very basic for editing.

1 SET UP SIGIL

Start by heading over to https://sigil-ebook.com/sigil/download, where you'll find download links for Windows (64-bit only, requires Windows 10 1809 or later), Mac (macOS 11 Big Sur or later), and Linux (available via your distro's repository, or install via flatpak for the latest build). After installing—click 'More Info > Run anyway' when prompted by the Smartscreen Filter—launch the app to reveal a three-paned window, as outlined in [**Image A**].

» In the left-hand pane is the Book Browser. This reveals the structure of an .epub file, which is in fact a zip file with an .epub file extension. Inside the file is everything required for your eBook, organized into various self-explanatory folders: Text contains your book text (formatted in XHTML format), while various resources from CSS stylesheets to physical images, fonts, and even embedded audio and video clips reside in their own folders.

» This emphasizes how Sigil isn't really suited as a tool for creating your book from scratch—its own editor requires you to work with the underlying XHTML code and rely on the Preview pane to check your edits as you go. Most people prefer working with a WYSIWYG editor, but before we get to that, let's look at some basic concepts about formatting books for reading on Kindles and other devices.

2 GET YOUR TEXT READY

Most people will have written their story in another app—Microsoft Word is the obvious choice here—but it's important to understand that whatever formatting you may have introduced in Word will have to be discarded or amended to suit the ePub format.

» First, most popular eBook formats—including the native ePub format favored by Sigil and Kobo, as well as Amazon's AZW3 format for its Kindles—are written in XHTML format. This means that your eBook uses basic header (<h1>, <h2> and so on) tags for structure and simple formatting. When it comes to choosing the default font and size, you can either leave it to the eReader's default settings (as chosen by its owner) or make use of ePub's support for CSS (Cascading Style Sheets) to not only choose a font for your entire document, but set up complex formatting, such as dedicated styles for different types of information you wish to display (such as callouts, block quotes, or side panels). However, for most people's needs, the basic header formats, along with simple styles (bold, italics, and so on) should be sufficient.

» When you come to edit your story, think about structure—this might be as simple as ensuring you have the work split into several chapters, or you may want a more complex structure comprising anything from volumes through chapters to sections and sub-sections. Why is this important? By styling these consistently using the header tags, you'll be able to automatically create a table of contents at the beginning of your book to allow readers to move between sections easily.

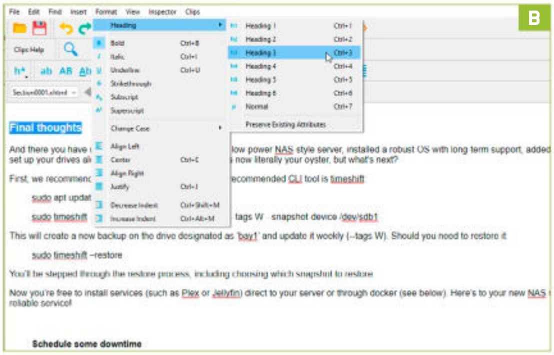
B
File Edit Find Insert Format View Inspector Clips
Clips Help
Section0001.xhtml
Heading
Bold Ctrl+B
Italic Ctrl+I
Underline Ctrl+U
Strikethrough
Subscript
Superscript
Change Case
Align Left
Center Ctrl+C
Align Right
Justify Ctrl+J
Decrease Indent Ctrl+Shift+M
Increase Indent Ctrl+Alt+M
Heading 1 Ctrl+1
Heading 2 Ctrl+2
Heading 3 Ctrl+3
Heading 4 Ctrl+4
Heading 5 Ctrl+5
Heading 6 Ctrl+6
Normal Ctrl+7
Preserve Existing Attributes

Final thoughts

And there you have … low power NAS style server, installed a robust OS with long term support, added … set up your drives al… now literally your oyster, but what's next?

First, we recommend… recommended CLI tool is timeshift

sudo apt updat…

sudo timeshift … tags W snapshot device /dev/sdb1

This will create a new backup on the drive designated as 'bay1' and update it weekly (--tags W). Should you need to restore it

sudo timeshift –restore

You'll be stepped through the restore process, including choosing which snapshot to restore

Now you're free to install services (such as Plex or Jellyfin) direct to your server or through docker (see below). Here's to your new NAS … reliable service!

Schedule some downtime

©SIGIL-EBOOK.COM

HOW TO NAVIGATE SIGIL

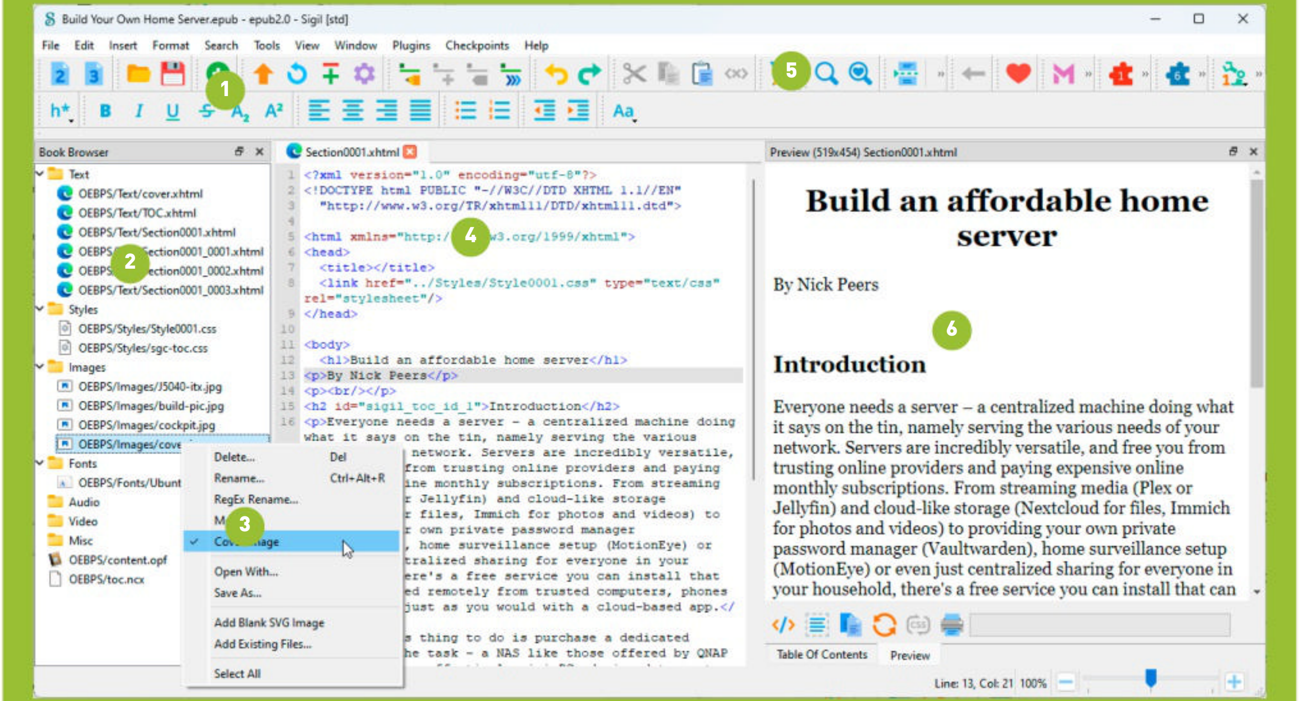
Build Your Own Home Server.epub - epub2.0 - Sigil [std]
File Edit Insert Format Search Tools View Window Plugins Checkpoints Help
1
2
3
4
5
6
Book Browser
Text
OEBPS/Text/cover.xhtml
OEBPS/Text/TOC.xhtml
OEBPS/Text/Section0001.xhtml
OEBPS/Text/Section0001_0001.xhtml
OEBPS/Text/Section0001_0002.xhtml
OEBPS/Text/Section0001_0003.xhtml
Styles
OEBPS/Styles/Style0001.css
OEBPS/Styles/sgc-toc.css
Images
OEBPS/Images/J5040-itx.jpg
OEBPS/Images/build-pic.jpg
OEBPS/Images/cockpit.jpg
OEBPS/Images/cove
Fonts
OEBPS/Fonts/Ubunt
Audio
Video
Misc
OEBPS/content.opf
OEBPS/toc.ncx
Delete... Del
Rename... Ctrl+Alt+R
RegEx Rename...
Cover Image
Open With...
Save As...
Add Blank SVG Image
Add Existing Files...
Select All
Section0001.xhtml
Preview (519x454) Section0001.xhtml
Table Of Contents
Preview
Line: 13, Col: 21 100%

```
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.1//EN"
  "http://www.w3.org/TR/xhtml11/DTD/xhtml11.dtd">

<html xmlns="http:/      w3.org/1999/xhtml">
<head>
  <title></title>
  <link href="../Styles/Style0001.css" type="text/css"
rel="stylesheet"/>
</head>

<body>
  <h1>Build an affordable home server</h1>
<p>By Nick Peers</p>
<p><br/></p>
<h2 id="sigil_toc_id_1">Introduction</h2>
<p>Everyone needs a server - a centralized machine doing
what it says on the tin, namely serving the various
             network. Servers are incredibly versatile,
            from trusting online providers and paying
            ine monthly subscriptions. From streaming
            r Jellyfin) and cloud-like storage
            r files, Immich for photos and videos) to
            r own private password manager
            , home surveillance setup (MotionEye) or
            tralized sharing for everyone in your
            ere's a free service you can install that
            ed remotely from trusted computers, phones
            just as you would with a cloud-based app.</

            s thing to do is purchase a dedicated
            he task - a NAS like those offered by QNAP
```

Build an affordable home server

By Nick Peers

Introduction

Everyone needs a server – a centralized machine doing what it says on the tin, namely serving the various needs of your network. Servers are incredibly versatile, and free you from trusting online providers and paying expensive online monthly subscriptions. From streaming media (Plex or Jellyfin) and cloud-like storage (Nextcloud for files, Immich for photos and videos) to providing your own private password manager (Vaultwarden), home surveillance setup (MotionEye) or even just centralized sharing for everyone in your household, there's a free service you can install that can

1. QUICK ACCESS TOOLBAR
As with all good applications, Sigil places some of its most popular tools on this toolbar—roll your mouse over one for a pop-up description.

2. BOOK BROWSER
The contents of your EPUB file, including text, images, CSS stylesheets, and other files (including fonts) are listed here.

3. ACCESS EPUB ELEMENTS
Double-click an item in the Book Browser to open it in the middle pane, or right-click to reveal more context-sensitive items depending on the item type.

4. CODE EDITOR
You can view the underlying XHTML or CSS that makes up your pages and stylesheets in this page, plus edit directly from here if you're comfortable with coding.

5. OPEN PAGE EDIT
Click this button to open your choice of third-party XHTML editor—in this tutorial we chose Sigil's companion app Page Edit, which offers WYSIWYG editing.

6. PREVIEW WINDOW
A preview of the currently selected file in the code editor is displayed here, updated in real time as you make changes.

3 FORMAT YOUR CONTENT

The best tool for formatting your content is Sigil's companion WYSIWYG editor, PageEdit. Download and install it from https://sigil-ebook.com/pageedit, following the same procedure you did for Sigil itself. Once installed, navigate to 'Edit > Preferences > General Settings' in Sigil, then click Browse under 'Set your preferred external xhtml editor', and navigate to C:\Users\{yourname}\AppData\Local\Programs\PageEdit to select PageEdit.exe.

» Once done, make sure Section0001.xhtml is displayed in the middle pane, and click the 'Launch External Xhtml Editor' button on the toolbar. This will open the file in PageEdit, where you'll see a blank word-processor-like page before you. Switch to Word, select all the text, copy it to the clipboard, then in PageEdit choose 'Edit > Paste', selecting Yes when prompted.

» Now, it's time to go through the content, styling it up using the 'Format > Heading' menu. Think carefully about the structure and hierarchy—for example, you may want to assign Heading 1 to the book title at the front, then Heading 2 for chapter titles, and Heading 3 to sub-section titles. That leaves three styles spare (Heading 4 to 6) that could be repurposed for other page elements, such as image captions or pulled out text or quotes, as shown in [Image B]. If you're willing to dip your toe into the murky waters of CSS to meet more complex formatting needs, check out the final box.

» Don't forget to press Ctrl + S regularly as you go—you'll see Sigil's main preview update to show your changes on-screen. You should also press Ctrl + S in Sigil itself to periodically save the parent ePub file, too.

4 ADD LINKS AND IMAGES

Once the headings are in place, go back through your copy, styling up the body text with bold, italics, and other supported elements (again, apply these through the Format menu). If you need to provide any weblinks, select the text that will house the hyperlink and choose Insert > Link, then type or paste the web address into the Target: box, before clicking OK. You can also create links to other parts of your book—you'll need to first place an anchor at the point you wish to jump to; to do this, position the cursor accordingly and then choose Insert > ID... before giving it a suitably descriptive code or name. Once done, you'll be able to select this ID from the list as shown in [Image C]. Again, click OK.

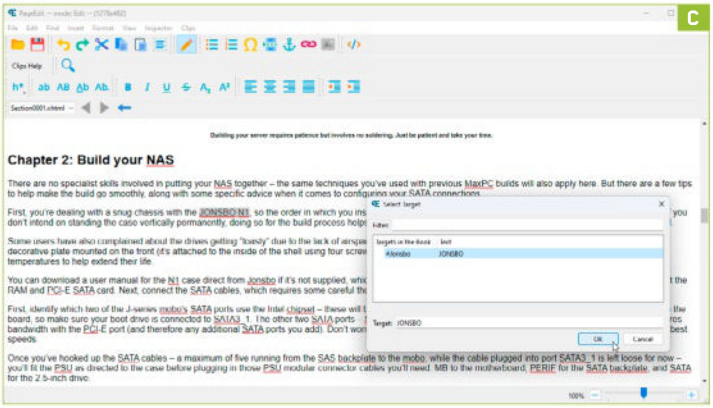

Building your server requires patience but involves no soldering. Just be patient and take your time.

Chapter 2: Build your NAS

There are no specialist skills involved in putting your NAS together – the same techniques you've used with previous MaxPC builds will also apply here. But there are a few tips to help make the build go smoothly, along with some specific advice when it comes to configuring your SATA connections.

Targets in the Book: #Jonsbo — JONSBO

Target: JONSBO

Once you've hooked up the SATA cables – a maximum of five running from the SAS backplate to the mobo, while the cable plugged into port SATA3_1 is left loose for now – you'll fit the PSU as directed to the case before plugging in those PSU modular connector cables you'll need: MB to the motherboard, PERIF for the SATA backplate, and SATA for the 2.5-inch drive.

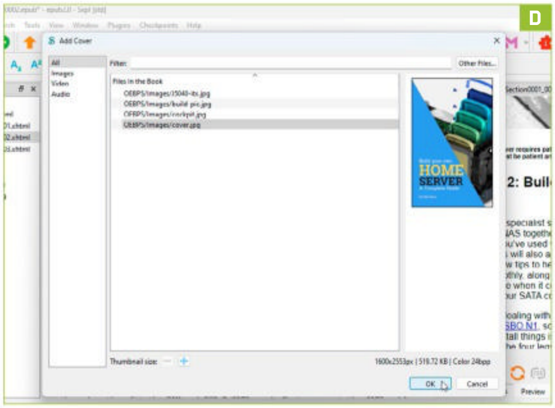

» If your book contains images (see the first box for tips on prepping them for use), you will need to embed these in your ePub file before you can insert them. Close the main PageEdit window and switch back to the main Sigil interface, right-click the Images folder in the left-hand pane, and choose 'Add Existing Files…' to browse for and add them. You can do the same with other files using the Audio, Video, and Misc folders, too.

» Now, position the cursor within the text in Sigil's code editor pane where you'd like the image to go, and choose Insert > File. Select your image from the list, and click OK. You should see something like the following appear in the code:

```
<img alt="build-pic" src="../Images/<image-name>.jpg"/>
```

» Now, reopen PageEdit from the shortcut in Sigil, where the image should now be visible, allowing you to align it centrally if you wish (note, while it's technically possible to insert images directly from PageEdit, it appears to be broken on our test machine).

5 SPLIT SECTIONS

Before proceeding further, look at the Checkpoints menu. This is an internal backup tool that lets you create versioned backups as you proceed, allowing you to revert to an earlier revision should you make changes you're subsequently unhappy with. While Sigil does have undo support, this failsafe should be triggered before applying any major changes to your book. To help identify each checkpoint, choose Checkpoints > Edit Checkpoint Descriptions. You can compare your current version with a previous checkpoint, too.

» As things stand, your eBook is one long continuous piece of prose. It makes sense to introduce proper breaks where you can—the obvious place is at the end of each chapter, which means readers start the next chapter on a fresh page. This is done in Sigil by splitting the single XHTML file into multiple files, one per chapter or section. It's easy to do, too—either from within Sigil's own code editor or PageEdit's friendlier interface. In either event, place the cursor where you want the first break to occur, and choose Insert > Split Marker.

» This inserts the following line of code into your book at the specified point:

```
<hr class="sigil_split_marker"/>
```

» Once you've placed markers where you would like them to occur, close PageEdit (if necessary) and choose 'Edit > Split At Markers' in Sigil. You'll immediately see that the single XHTML file has now been split into multiple files, listed separately in the Text folder.

6 ADD A COVER

The most important image in your eBook will, of course, be the cover. You can do this by hand from scratch using any image editor or design tool capable of outputting files to JPEG—look to see if any eBook templates are supplied, and if not, create a blank file 2,560 pixels high by 1,600 pixels wide as per Amazon's recommendations.

PREPPING YOUR IMAGES

If you'd like to include images within your eBook, then Sigil has the tools you need. But before you add them to your document, you need to make sure they fit certain criteria: first, they'll need to be in GIF or JPG format (PNG may work, but it's not guaranteed)—JPG is best for keeping the file size down. Second, make sure their DPI is set between 72 and 150 dpi (you should find an option for this under your image editor's resize dialog), and that they're no bigger than 1,000 pixels on their longest side.

We also recommend converting the image to grayscale, as most eReaders can't display color, and it'll help reduce the file size further, which could have a major impact if you've got lots of images to include in your book.

You can do all this by hand, but if you have dozens of image files, try a batch conversion tool like Faststone Photo Resizer (www.faststone.org). Load your files, choose your output format and directory, and click Advanced Options to access various tabs to make your adjustments en masse.

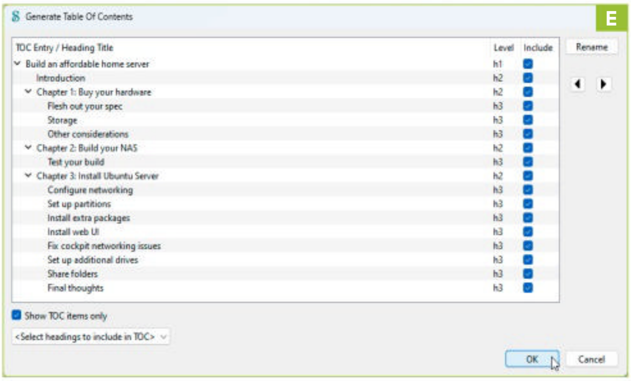

TOC Entry / Heading Title	Level	Include
Build an affordable home server	h1	☑
Introduction	h2	☑
Chapter 1: Buy your hardware	h2	☑
Flesh out your spec	h3	☑
Storage	h3	☑
Other considerations	h3	☑
Chapter 2: Build your NAS	h2	☑
Test your build	h3	☑
Chapter 3: Install Ubuntu Server	h2	☑
Configure networking	h3	☑
Set up partitions	h3	☑
Install extra packages	h3	☑
Install web UI	h3	☑
Fix cockpit networking issues	h3	☑
Set up additional drives	h3	☑
Share folders	h3	☑
Final thoughts	h3	☑

☑ Show TOC items only

<Select headings to include in TOC>

©SIGIL-EBOOK.COM

» If you're unable to find inspiration from a blank canvas, try a free online eBook cover design tool. Snappa (https://snappa.com/templates/ebook-cover) is one of the few online sites to allow you to both create and download—up to three files a month—your finished eBook cover for free. Once signed up for your free account, click eBook Cover under 'Blogging and Infographics' to design your cover using one of well over 100 templates. Once complete, save it as a web-optimized JPG for embedding into your book.

» Once done, make sure the file is named cover.jpg, then import it into Sigil as a regular image. Once done, choose Tools > Add Cover, select cover.jpg, as shown in [**Image D**], and click OK. Sigil will create another page (cover.xhtml) whose sole purpose is to display the cover at the front of your book. If you right-click the cover.jpg file under Images, you should see that 'Cover Image' is checked, confirming it's correctly set up.

7 ADD BOOK DETAILS AND CREATE TOC

Next, reopen the Tools menu, but this time select Metadata Editor... Fill the existing Title tag with the name of your book, then choose Add Metadata to add more information using the list of available tags. These include Author, Description, Date, and various identifiers, including ISBN number, should you already have set this up or previously published the book.

» Once done, click OK. Assuming you've structured your book using headings, choose 'Tools > Table of Contents > Generate Table of Contents'. A window will open showing all existing headings in your document in the order in which they appear.

» You'll see that by default all headings are selected, which can lead to a messy ToC if you've used some headings for non-header styles (such as captions or box text). Thankfully, you can uncheck unwanted items or—better still—click the <Select headings to include in TOC> dropdown to restrict them to a certain level (say level 2 or level 3). Once selected, only those headings meeting your chosen criteria will be shown, as can be seen in [**Image E**], and you can still remove individual headings using the checkboxes as before. When you're done, click OK.

» As things stand, the table of contents is internal only—to add a TOC to the actual book, choose 'Tools > Table of Contents > Create HTML from Table of Contents' to create a new TOX.xhtml document. If the file appears above cover.xhtml under Text, drag it beneath it so it appears in the correct order of your book.

8 VALIDATE, EXPORT AND TEST

Once your book is completed to your satisfaction, perform some quick checks before viewing it in an eReader. First, choose Edit > Preferences > Spellcheck Dictionaries to verify that the correct language is selected (English—United States should already be selected), then select Tools > Spellcheck > Spellcheck to generate a list of words not recognized by the dictionary. Then go through the list using the controls on the right to ignore, and add to dictionary or provide a correction to the misspelled word as required.

» Once you're happy, click OK and choose Tools > Well-Formed Check EPUB to bring up a separate panel that will list any potential problems with the file. Other tools worth exploring in conjunction with the user guide include reformatting the underlying XHTML and validating your CSS stylesheets if applicable.

» When your book has passed these tests, try opening the EPUB file in a dedicated EPUB viewer like Sumatra PDF Reader (www.sumatrapdfreader.org/free-pdf-reader) to see how it looks—as shown in [**Image F**]. You can also transfer EPUB files directly to a eReader for further tests—use Calibre (https://calibre-ebook.com) to transfer it to a non-Kindle reader, or Kindle Previewer (https://kdp.amazon.com/en_US/help/topic/G202131170), which will automatically convert it to AZW3 for reading on a Kindle. ⏻

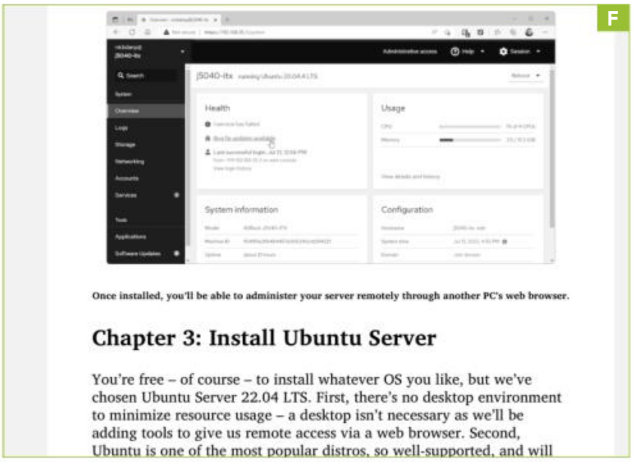

F

Once installed, you'll be able to administer your server remotely through another PC's web browser.

Chapter 3: Install Ubuntu Server

You're free – of course – to install whatever OS you like, but we've chosen Ubuntu Server 22.04 LTS. First, there's no desktop environment to minimize resource usage – a desktop isn't necessary as we'll be adding tools to give us remote access via a web browser. Second, Ubuntu is one of the most popular distros, so well-supported, and will

STYLE YOUR BOOKS WITH CSS

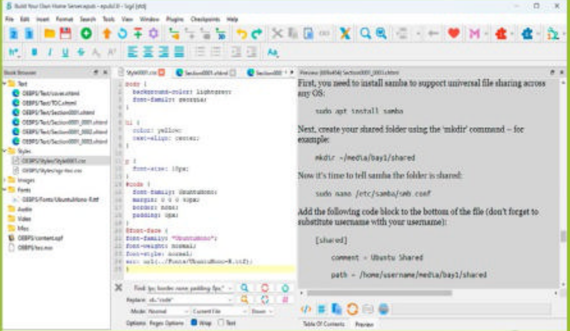

First, you need to install samba to support universal file sharing across any OS:

```
sudo apt install samba
```

Next, create your shared folder using the 'mkdir' command – for example:

```
mkdir ~/media/bay1/shared
```

Now it's time to tell samba the folder is shared:

```
sudo nano /etc/samba/smb.conf
```

Add the following code block to the bottom of the file (don't forget to substitute username with your username):

```
[shared]
   comment = Ubuntu Shared
   path = /home/username/media/bay1/shared
```

CSS is used to style ePub documents beyond the basics—CSS allows you to define everything from the typeface used along with its style, size, color, and background, to the layout of individual elements on the page. These can be universal settings or limited to both existing styles and brand new ones created from scratch.

It requires learning how CSS works and what you can do with it—see www.w3schools.com/css for an excellent introduction to the subject. Sigil supports multiple CSS files (one's created for your table of contents, for example), so to create a CSS file for the rest of your document, right-click the Styles folder and choose 'Add Blank Stylesheet'. Next, select all the XHTML files under Text that you want the stylesheet to apply to, before right-clicking and choosing 'Link Stylesheets' to select your new stylesheet.

You can enter CSS code directly into Sigil—we recommend using W3 Schools's online editor to build your code, then copy and paste everything inside the <style></style> tags into the CSS file in Sigil (you'll see the preview of any linked XHTML file updates as the code goes in), which you can then tweak to perfection.

Make secure backups using an ISO file

YOU'LL NEED THIS

WINCDEMU
https://wincdemu.sysprogs.org

YOU'VE ALMOST CERTAINLY come across an ISO file before. They're typically used to bundle huge amounts of data, such as in Windows installers, in a widely compatible format. With the format being so predictable, it's not surprising that new uses for it have since been found. One is to create virtual optical discs, which exist purely as bundles of data, and have no physical presence at all. The data in the bundle complies with the ISO standards, so it can be opened reliably and repeatedly across multiple operating systems, which makes it a safe format for producing manual backups. Because reading them doesn't rely on a specific piece of software, you should be able to open an ISO file for years to come. –**NIK RAWLINSON**

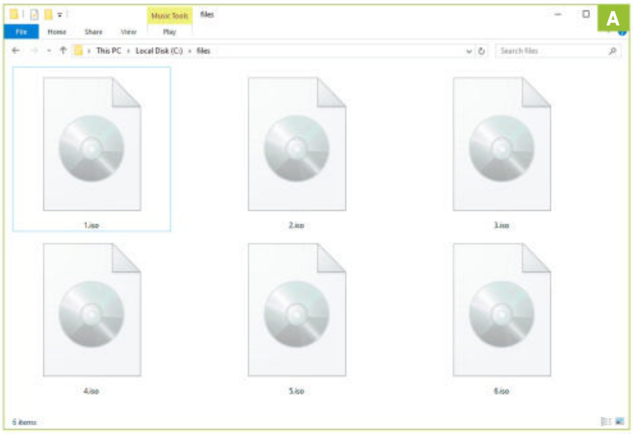

1 ISO AS PART OF YOUR BACKUP STRATEGY

We've mentioned the 3-2-1 rule of backup before: always maintain at least three complete copies of your data across two different media, of which one should be off-site (that is, not in your home). If you're going to follow our advice here and you're not also using an online backup service to store your data off-site, we recommend combining it with backups made to external hard drives or USB sticks using regular backup software.

» You should also ask a friend or family member to look after one of the copies made using the backup software. Why that and not an ISO file? Because most backup software lets you encrypt and password-protect your backups. We don't want to do that with our ISO file because we want it to be easily accessible for years into the future. If it's encrypted using a standard that falls out of favor or we forget the password, then that won't be possible.

» As we'll explain, the ISO backup of your data will behave just like a regular unencrypted drive. So, if you were to hand it over to someone else and they opened it on their machine, it would appear in the File Explorer sidebar like any regular external drive. They'd then be able to navigate through your files, read your personal documents, look at your photos, and so on. [**Image A**]

2 INSTALL WINCDEMU

While Windows can open (or 'mount') an ISO file natively—meaning without using any third-party software—it can't create them without using another program. There are plenty of options, but we're going to use the free software WinCDEmu.

» The reason we've opted for this over its rivals is that it adds a right-click shortcut to File Explorer for converting folders into ISOs. Because backing up is something you should do on a regular basis, this could prove to be a time-saver. Visit wincdemu.sysprogs.org, then click the green Download button. The installer should download in a few seconds. Once it has, open the 'WinCDEmu-4.1.exe' file from your Downloads folder.

» Windows will ask if you're sure you want the program to make changes to your system—click Yes. By default, WinCDEmu will be installed to a subfolder of the 'Program Files (x86)' folder [**Image B**]. If you're happy with this, click Install. If not, tick the box beside 'Customize installation options', then choose a different location before clicking Install. As part of the process, Windows Security will pop up an alert asking if you want to install device software. Again, click Install to proceed.

3 CREATE YOUR FIRST ISO FILE

Create an ISO of your Documents folder as a test. Open File Explorer and click the C: drive in the sidebar. If this isn't visible, scroll to This PC and check the triangle to the left. If it's pointing right, like a chevron (>), click it to reveal your attached drives, then click C:.

» Now, navigate through your file system by double-clicking Users, then your username, until you reach your

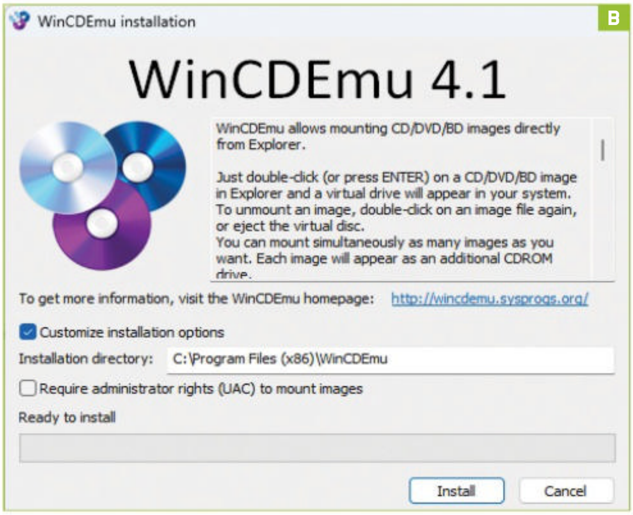

WinCDEmu allows mounting CD/DVD/BD images directly from Explorer.

Just double-click (or press ENTER) on a CD/DVD/BD image in Explorer and a virtual drive will appear in your system. To unmount an image, double-click on an image file again, or eject the virtual disc.
You can mount simultaneously as many images as you want. Each image will appear as an additional CDROM drive.

To get more information, visit the WinCDEmu homepage: http://wincdemu.sysprogs.org/

☑ Customize installation options

Installation directory: C:\Program Files (x86)\WinCDEmu

☐ Require administrator rights (UAC) to mount images

Ready to install

© WINCDEMU

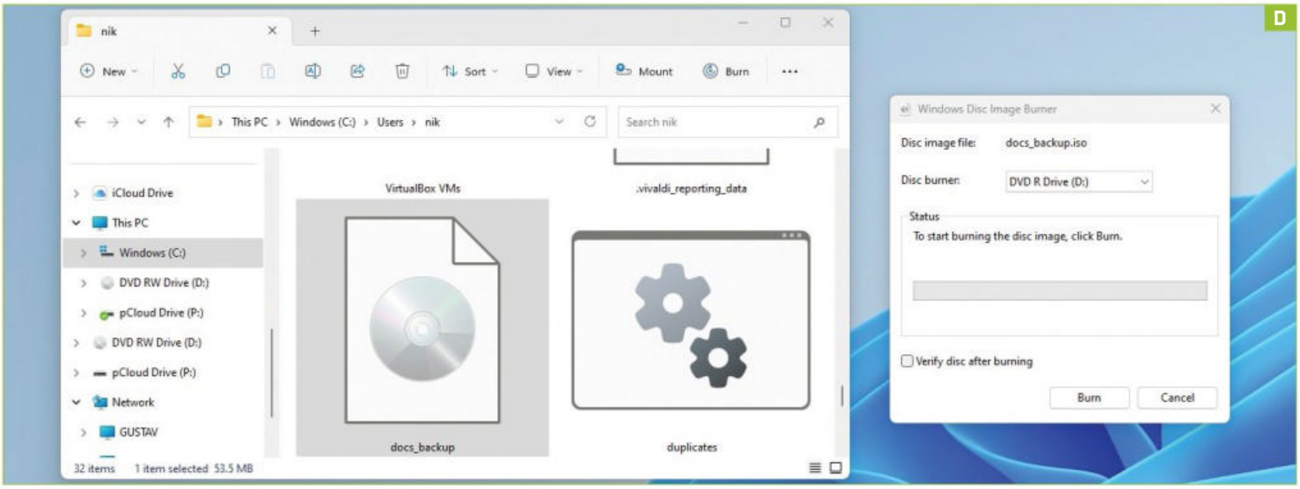
D
nik
New
Sort
View
Mount
Burn
This PC › Windows (C:) › Users › nik
Search nik
iCloud Drive
This PC
Windows (C:)
DVD RW Drive (D:)
pCloud Drive (P:)
DVD RW Drive (D:)
pCloud Drive (P:)
Network
GUSTAV
VirtualBox VMs
.vivaldi_reporting_data
docs_backup
duplicates
32 items 1 item selected 53.5 MB
Windows Disc Image Burner
Burn
Cancel

Disc image file: docs_backup.iso

Disc burner: DVD R Drive (D:)

Status
To start burning the disc image, click Burn.

☐ Verify disc after burning

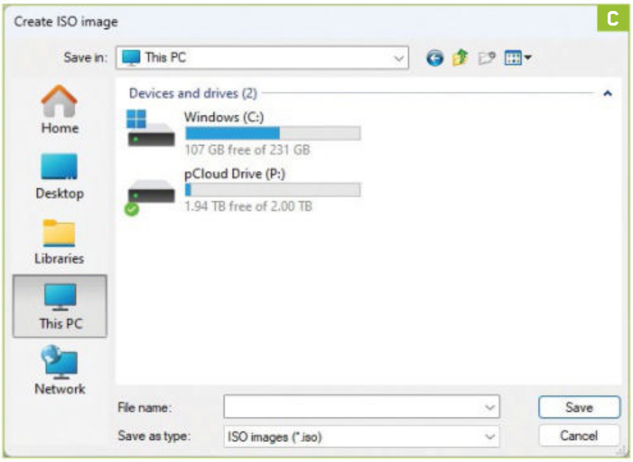
Create ISO image
C
Home
Desktop
Libraries
This PC
Network
Devices and drives (2)
Windows (C:)
107 GB free of 231 GB
pCloud Drive (P:)
1.94 TB free of 2.00 TB
Save
Cancel

Save in: This PC

File name:

Save as type: ISO images (*.iso)

user folder. Right-click the Documents folder and, in Windows 10, click 'Build an ISO image'. If you're using Windows 11, you'll need to perform one more step: after right-clicking Documents, click 'Show more options' to find the 'Build an ISO image' option.

» Navigate to the location where you'd like to create your ISO file. Although WinCDEmu lets you save the file to the same drive as the one with your source files, you should avoid doing this if you're creating a backup. It not only wastes space by duplicating files on that drive; it doesn't protect against your hard drive failing. Instead, navigate to an external drive or, if you plan to write the ISO file to a DVD or CD, store it temporarily on your internal drive before copying and deleting this temporary ISO.

» Your drives might not be visible immediately, in which case, click This PC [**Image C**] in the 'Create ISO image' box to reveal them, then click the one you want. Name the file, leave the 'Save as type' menu set to the default—'ISO images (*.iso)—then click Save. WinCDEmu will display a progress bar showing how far it has got creating your backup. When complete, click OK, followed by Close.

4 TEST YOUR ISO FILE

It's essential that you test your ISO file to make sure it's both accessible and an accurate copy of the files you want to secure. This is easy, because Windows can open your ISO like an attached drive.

» Open File Explorer and navigate to the drive and folder in which you saved the ISO file. Double-click the file, and it will mount in the sidebar. Our ISO has been given the drive letter D, and is being treated as a CD-ROM, even though we don't have an internal optical drive in the mini PC we're using to write this feature. Click through some of the folders and open a selection of files in the ISO to make sure they behave as expected.

» Keep the ISO somewhere safe so you can restore your files from it if disaster strikes, and don't forget to create a replacement every few weeks. If you want to write the ISO to a disc, eject the ISO file by right-clicking it in the File Explorer sidebar and clicking Eject. Now, insert a writable DVD or CD into your optical drive and, when Windows has detected the blank disc, select your ISO file in File Explorer. When you do, a new Burn button will appear on the top toolbar [**Image D**].

» We've attached an external DVD burner to our PC, which has taken the D: drive name because ejecting the ISO file freed up that slot. Make sure your burner is selected in the 'Disc burner' menu, then click Burn. ⏻

HOW TO USE ISO FILES

I most often encounter ISO files when I download new software or operating systems from the web. Windows arrives this way, but fortunately I rarely have to reinstall it from scratch. Mostly, I use ISO files for operating systems on single-board computers like Raspberry Pi, which I use for a range of home monitoring, weather, and information tasks. That said, writing this has prompted me to use ISO files to back up large blocks of work, including novels and non-fiction books.

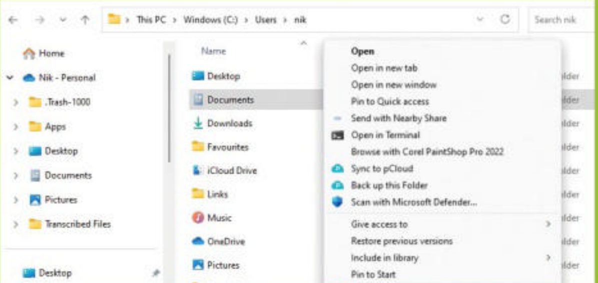
This PC › Windows (C:) › Users › nik
Search nik
Home
Nik - Personal
.Trash-1000
Apps
Desktop
Documents
Pictures
Transcribed Files
Desktop
Name
Desktop
Documents
Downloads
Favourites
iCloud Drive
Links
Music
OneDrive
Pictures
Saved Games
Open
Open in new tab
Open in new window
Pin to Quick access
Send with Nearby Share
Open in Terminal
Browse with Corel PaintShop Pro 2022
Sync to pCloud
Back up this Folder
Scan with Microsoft Defender...
Give access to
Restore previous versions
Include in library
Pin to Start
Build an ISO image

LAB NOTES

JEREMY LAIRD, CONTRIBUTOR

Is OLED burn-in actually a thing?

Could OLED monitors be going the way of early SSDs?

HOW MUCH OF A THING is OLED burn in? That's hard to answer in normal testing. We typically get a few weeks at best to review something like a gaming monitor. That's plenty for assessing most aspects of a display. You can get a good feel for the image quality, build, features, and so on. But burn in? That's much harder to judge.

Of course, long-term durability is always difficult to measure. I remember when the first consumer SSDs came out to fairly universal rave reviews. Then it turned out the drives wore out pretty quickly. While SSD durability remains difficult to gauge, the technology has improved. But what about OLED panels in PC monitors?

If you look at the latest long-term OLED test results from Rtings, just about the only outfit currently having a stab at this problem, they make for worrying reading—for monitors with Samsung's QD-OLED panels, at least. Rtings only has three OLED monitors in their test. But after around 3,000 hours of testing, both of the monitors with Samsung panels show signs of burn in. The third monitor has an LG WOLED panel, and as far as I can tell, seems to be fine.

The Rtings test routine is the worst-case scenario. The question is to what extent burn in will be a problem in more typical PC usage—and that is very difficult to answer.

Then there's the Windows interface. The taskbar seems like a prime candidate for burn in, as do desktop icons. It's likely possible to mitigate that through various measures, such as setting Windows to dark mode and the taskbar to autohide. But at some point, you have to ask yourself if this technology is fit for purpose.

I've been bullish about OLED durability. I haven't seen any issues myself, but I have to admit that I'm a bit spooked by the Rtings test results. I think it's going to be a while before we have a really good feel for how well OLED truly stands up to PC usage.

WITCHER
WILD HUNT
ALIENWARE

Alienware's monitors use those potentially problematic Samsung QD-OLED panels.

SAM LEWIS

Staff Writer

Grown-up fans of gaming have a problem. Overtly game-centric seating can look silly in an office setting. Some would say that's fine—the office is no place for games. But realistically, work settings are pretty fluid for a lot of us. An office chair that's also good for gaming? Yes please.

Enter the NeueChair from Secretlab. At $729, it's not going to appeal to everyone. But if you can afford it, you're getting one heck of a chair. With its mesh backing and adjustable lumbar support, it's not squishy, but does encourage good posture. If you must recline, it will do that, but only to 130 degrees less than the virtually horizontal poses that many gaming chairs can achieve.

The NeueChair is built to last. That's backed up by a hefty 12-year warranty, which goes some way to offsetting that painful price.

Only the armrests disappoint. You can raise them on a ratchet mechanism, but they don't go high enough for my desk. They also don't move to support your elbows, and you don't get a headrest, which adds to the cost of an already expensive chair. But since this is more suited to the active position than reclining, it's not a major loss.

What matters is that this is the most comfortable office chair I've ever sat in. If you're going to be at your desk most of the day working, and into the night gaming, then it's worth the investment.

© ALIENWARE

Editor's Pick:
Framework 13 AMD mainboard

The ultimate hybrid gaming ultrabook

THE NEW Framework AMD mainboard is here. I already had a crush on Framework's laptops and the notion of a system that can be upgraded down the line. But while Frameworks have been great work machines, this hasn't been the case for gaming. But dropping the new AMD Ryzen 7 7840U board into this old chassis? That's a game changer.

My first taste of the Framework 13 was courtesy of an 11th Gen Intel Tiger Lake CPU. The 13-inch scale is great for a predominantly office-focused notebook, the design is slim and simple, the modular ports are great, and the 16:10 aspect ratio of the sharp, glossy 2,256 by 1,504 resolution screen is just great.

Even after adding a newer 12th Gen Alder Lake chip, and then a 13th Gen Raptor Lake board, those updates only made it a better office machine. The new AMD Zen 4 APU upgrade doesn't make this a gaming beast, but the Ryzen 7 7840U might just be the best chip AMD makes.

It's finding its way into all those new gaming handhelds, including the Asus ROG Ally. With 12 RDNA 3 graphics compute units, the 7840U has game in a way no previous integrated GPU has. As with those handhelds, that GPU translates into truly usable 1080p gaming performance.

What isn't included in many games is AMD's new Frame Generation technology that sits alongside its FSR 3 rollout. This is a shame—when it works, you're looking at a boost in frame rates from enabling both FSR 3 and AMD's Frame Gen feature.

Still, the hardware is effective at native settings. However, it's not quite as quick in gaming terms as the handhelds with the same APU inside. It's often only a few frames per second behind the likes of the ROG Ally, but those dedicated gaming devices do seem to have an edge.

Compare it with the Core i7 1370P, and you're talking about a 14-core chip, with 20 processing threads beaten by this eight-core Zen 4 chip. The single-thread Cinebench R23 performance of the Intel Raptor Lake CPU is better, but that's where Intel's advantage ends. In multi-threaded tasks, where you'd expect the higher Intel thread count to win, it's the AMD chip that pulls ahead, whether that's in Cinebench, Blender rendering, or HD video encoding.

You might think that this extra gaming performance and processing power would demand more juice from the battery at the heart of the Framework 13, but the opposite is true. The AMD board is actually more efficient. Using the upgraded 61Wh battery, the Intel board manages 68 minutes using PCMark's Gaming battery life test. The AMD mainboard? 86 minutes.

Oh, and did I mention that the AMD option is actually cheaper? The Core i7 1370P mainboard retails for $1,049 on its own—a hefty upgrade that's the price of an entire gaming laptop of decent spec. This AMD Ryzen 7 7840U mainboard is $699. Still not exactly cheap, but worth it. **–JL**

$699, www.frame.work

The new AMD mainboard for the Framework 13 is a genuine game changer.

© AMD, PODMAN, SEGA

Reviewed...

74 Sapphire Pulse RX 7700 XT

76 Samsung Odyssey Neo G9 G95NC

78 Ayaneo Air 1S

80 MSI Prestige 13 Evo

82 Asus Vivobook S 15 OLED

84 Asus ROG Zephyrus M16

87 Logitech Pro X Superlight 2

88 Phanteks NV7

89 Corsair K70 Max

podman
Podman is running
Version 4.7.1
Podman Machine

90 Total War: Pharaoh

92 Docker vs Podman

The Sapphire Pulse RX 7700 XT is overshadowed by AMD's own 7800 XT boards.

Sapphire Pulse RX 7700 XT

Pretty nice, but there's a problem with the price

WITH THE RELEASE of the RX 7800 XT and RX 7700 XT, AMD's RX 7000-series is apparently complete. These cards fill the gaping hole AMD had in its line-up following the release of the high-end RX 7900 XT and RX 7900 XTX in late 2022 and the mainstream RX 7600 in May.

Specifically, we're dealing with the Sapphire Pulse RX 7700 XT. All 7700 XTs are AIB cards, as there is no AMD reference card. This is the base model of the three cards Sapphire offers. It's joined by the mid-range Pure and high-end Nitro+. The Pulse is priced at the 7700 XT's $449 MSPR, and honestly, spending more on premium models is foolish.

Anyway, the 7700 XT looks good compared to its 6700 XT predecessor. It has got more shaders, better claimed ray-tracing performance, faster memory, and includes AI accelerators the RX 6000-series cards specifically lacked.

Nice, but the 7700 XT's problem is none other than the RX 7800 XT that costs only $50 more, but offers a hefty spec bump. Both use the same Navi 32 chiplet GPU. The 7700 XT packs 54 Compute Units, 54 Ray Accelerators, and 96 ROPs, while the 7800 XT clocks in with 60, 60, and 120.

The 7700 XT's core clock is 100 MHz higher, but it only has 12GB of 18Gbps GDDR6 memory over a 192-bit bus. The fully enabled RX 7800 XT rocks 16GB of 19.5 Gbps memory, a 256-bit bus, and more Infinity Cache. You also get DisplayPort 2.1, which no Nvidia GPU currently offers. On the other hand, Nvidia counters with DLSS 3 and more mature Frame Generation, not to mention ray reconstruction.

More broadly, AMD still lags far behind Nvidia when it comes to ray tracing and the 7700 XT is no exception. But for regular raster performance, the 7700 XT has both Nvidia's slightly cheaper $399 RTX 4060 Ti 8GB and slightly more expensive $499 RTX 4060 Ti 16GB cards comprehensively beaten. The RX 7700 XT is a capable 1440p card in a way that the 4060 Ti simply isn't.

The exception is *Cyberpunk 2077*, but that runs like plucked turkey on anything less than an RTX 4090. 4K in demanding titles remains the domain of more expensive cards. But with older games, competitive shooters, or MOBA games that are well optimized, the RX 7700 XT will run those at 4K without a hitch.

Likewise, the Sapphire Pulse cooler is capable of handling the 230W TBP of the card, and in our testing this card draws less power than a 6700 XT, though the 4060 Ti uses less power. Still, the Pulse remains remarkably quiet, even if the backplate does get rather hot.

This is a pretty nice card by many measures, then, but its major problem is pricing. Certainly, any premium priced 7700 XT is dead in the water. Why would you consider a 7700 XT that costs as much as a faster RX 7800 XT? Moreover, the price of the 7700 XT both feels too close to the 7800 XT and is bound to fall over.

In a month or six, there's a good chance that the 7700 XT will be priced nearer the $399 mark it arguably should have had from launch. It will then be a pretty nice 1440p option, and offer a big enough raw performance advantage over the 4060 Ti to compensate for AMD's weaker features in the scaling, frame generation, and ray-tracing departments.

It's also worth noting that at the time of writing, AMD is bundling *Starfield* with RX 7700 XT's. That alone might convince you to jump on the RX 7700 XT over the RTX 4060 Ti. But right now, at $449, this card doesn't make a lot of sense. You'd be much better off stretching another $50 for a 7800 XT. **–CHRIS SZEWCZYK**

VERDICT
7

Sapphire Pulse RX 7700 XT

+ PULSE RACING Gives the RTX 4060 Ti a nasty beating; Good 1440p performance.

− TOTALLY FLATLINED Priced too close to the 7800 XT; Weak ray tracing

$449, www.sapphiretech.com

BENCHMARKS

	Sapphire Pulse RX 7700 XT	RTX 4060 Ti 8GB
Cyberpunk 2077 RT 1080p (fps)	37	**44**
Cyberpunk 2077 RT 1440p (fps)	23	**26**
F1 2022 RT 1080p (fps)	**118**	106
F1 2022 RT 1440p (fps)	**58**	54
Red Dead Redemption 2 1080p (fps)	**108**	102
Red Dead Redemption 2 1440p (fps)	**88**	79
Total War: Warhammer III 1080p (fps)	**124**	105
Total War: Warhammer III 1440p (fps)	**79**	68

Test platform: Intel Core i9 12900K, Asus ROG Z690 Maximus Hero, 32GB G.Skill Trident Z5 RGB DDR5-5600, 1TB WD Black SN850

SPECIFICATIONS

Architecture	RDNA 3
Manufacturing process	5nm GCD + 6nm MCD
Compute units	54
Shaders	3456
ROPs	96
Infinity cache	48MB
Memory	12GB 18Gbps GDDR6
Memory bus	192-bit
TBP	245W

© SAPPHIRE

Samsung Odyssey Neo G9 G95NC

Samsung's new 57-inch dual-4K monster is so spectacular it's silly

IS THERE SUCH thing as too much screen? Obviously, yes. Try putting a 65-inch HDTV on your desktop, and the limitations of size are clear. But the new Samsung Odyssey Neo G9 G95NC Dual UHD is no thinly disguised TV; it's a PC-optimized panel. Yet, this still has us wondering if you can have too much of a good thing.

What you're looking at, in effect, is a pair of 32-inch 240Hz 4K gaming monitors fused into a single 57-inch ultra-curved panel. That's no fewer than 7,680 by 2,160 pixels, or just over 16.5 million. Ouch.

Dual 4K running at such a high refresh rate is unprecedented. Indeed, it's so novel that only AMD's latest RDNA 3-powered Radeon RX 7000 Series GPU can hit the full 240Hz, courtesy of their DisplayPort 2.1 interfaces. Nvidia's RTX 40-series GPUs are limited to DP 1.4, and can only do 120Hz. Of course, you could call that an academic distinction. Good luck hitting 240Hz at dual 4K in *Cyberpunk* with all the ray-tracing twangers maxed out, even on an RTX 4090.

For HDR support, you get DisplayHDR 1000 certification and 2,392 dimming zones. The panel itself is VA, and pixel response is rated at 1ms—quick for an LCD monitor, but miles behind the 0.1ms or even lower of the latest breed of OLED panels. As for connectivity, along with the DisplayPort socket, there's a trio of HDMI inputs, plus USB-C. However, the latter is only for the USB hub.

Out of the box, this is large. We're used to Samsung's 49-inch G9s, but somehow this 57-inch monster is still shocking. First impressions are of a monitor that comfortably exceeds the claimed brightness rating. Samsung quotes 420 nits typical for SDR content, but this is punchier than your usual 400 nit panel.

This dual-4K monster does things no other display can match.

HDR performance is beyond the HDR1000 certification, too. Pan the camera up to the sky and the desert sun in *Cyberpunk 2077*, and the panel sizzles. HDR video content looks spectacular, too.

There are limitations, of course. Even with a couple of thousand dimming zones, the backlight resolution is low. However, Samsung has sorted out how to handle SDR content in HDR mode. You have no backlighting weirdness with SDR content. Indeed, we didn't see any of the glitches and bugginess that have blighted Samsung's big Odyssey gaming LCDs.

What we did see is the elevated pixel density on an epic scale. The 140DPI is exactly the same as a 32-inch 4K monitor, it's just you've never seen this kind of pixel density on such a large panel before.

That said, it's questionable how much you notice this in-game. At normal game viewing distances, the heightened image detail is definitely there. It's just not super obvious versus Samsung's smaller and lower-resolution G9 monitors.

Likewise, it's debatable how much benefit the ultrawide 32:9 aspect is for gaming. Sure, it's fantastic for that wrap-around feel of immersion. But most first-person shooters don't support it properly in FoV terms, so you end up with lots of stretching and distortion at the edges.

The visual ergonomics in terms of UI elements with such a huge, ultrawide screen are problematic, too. Having to crane your neck just to glimpse your minimap or health status isn't exactly ergonomic. So, the funny thing about the Samsung Odyssey Neo G9 G95NC is that it's more appealing as a PC monitor than as a gaming panel. –**JEREMY LAIRD**

VERDICT

9

Samsung Odyssey Neo G9 G95NC

⊞ **GAME ON** More pixels than you can possibly imagine; Well-optimized local dimming.

⊟ **GAME OVER** So big and wide, it's actually unwieldy for gaming; Not exactly cheap.

$2,499, www.samsung.com

SPECIFICATIONS

Screen size	57-inch
Resolution	7,680 x 2,160
Brightness	420 nits SDR
Color coverage	95% DCI-P3
Response time	1ms
Refresh rate	240Hz
Contrast ratio	2,500:1
Features	VA panel, adaptive sync, 1x DisplayPort 2.1, 3 x HDMI 2.1, USB-C hub, 1000R curve

© UNSPLASH

Ayaneo Air 1S

A seemingly impossible handheld gaming PC

FOR THE AYANEO AIR 1S, the default comparison is obviously the Steam Deck, right? Size-wise, the Ayaneo Air 1S is actually more like a Nintendo Switch. In fact, it's a touch smaller than the Switch, or slightly bigger than a medium-large banana. It really is incredibly compact, this handheld wonder.

But here's the thing: it comes with the same incredible AMD Ryzen 7 7840U chip found in the AOKZOE A1 Pro or Asus ROG Ally. That's a full eight-core, 16-thread Zen 4 processor and Radeon 780M iGPU powered by 12 RDNA 3 CUs—four more CUs than the Steam Deck's RDNA 2 chip.

Then there's the Ayaneo's memory and SSD configuration. Our review model comes with 32GB of LPDDR5X memory and a 2TB 2280 NVMe SSD. Yeah, a full-length 2280 drive. If you're familiar with the Steam Deck, or really most PC gaming handhelds, you'll know that they use the compact 2230 SSD form factor. This thing is tiny, but it's still an absolute beast in pure spec terms.

With that inevitably comes an equally beastly price, namely $1,029 as specified here. A budget alternative to the Steam Deck this is not. But then for a 32GB, 2TB, eight-core Zen 4-powered PC, it's not an awful price, especially when you consider the extreme form factor and the other elements you get for the money.

That includes a 5.5-inch 1080p AMOLED screen. It's not the biggest panel for in-depth strategy games or tons of text, but it delivers a very crisp image. You also get a full suite of handheld gaming controls. The dual hall-effect analog sticks are great for gaming, though Steam Deck owners might miss that device's trackpads for Windows navigation. Plugging in a mouse would be better, and is easily done thanks to dual USB4 Type-C connections and Bluetooth.

There's also a desktop shortcut button, and two hidden buttons along the top of the device, one of which is set to promptly bring up the on-screen keyboard and the fingerprint scanner for Windows Hello. All of these help with the Windows 11 experience, which is pretty polished.

Indeed, the OS feels as convenient as it gets. Fans of the Steam Deck's SteamOS, powered by Linux, might be surprised by how well Windows 11 suits the compact Air 1S. Part of that is the pop-up control panel, reminiscent of the Steam Deck's, which can be brought up at the push of a button to control the Ayaneo at a system level. Overall, the Ayaneo gives you a lot of what's great about the Steam Deck's OS in a device that can also run *Destiny 2*, *Fortnite*, or any other game that has been forbidden from Valve's open OS.

That said, this compact device is still better suited to indie darling than triple-A gaming. Its top TDP is a modest 25W, but on battery power that drops to just 20W. Some handhelds run the Ryzen 7 7840U chip at up to 30W, and that power deficit does limit performance. For some, that may be a deal breaker. Certainly, there are faster alternatives. Likewise, if the Ayaneo isn't going to major on performance, why pass up on the Steam Deck's established ecosystem? Oh, or the Deck's battery life. Even in 20W mode, the Ayaneo only lasts a little over an hour in PCMark 10's gaming benchmark. Ouch.

So the question to ask yourself is if the compact size is worth it over the likes of the Steam Deck and the Asus Ally. For the sheer convenience of a handheld gaming PC that you don't have to think twice about packing into almost any bag or backpack, the Ayaneo Air 1S is super appealing. Yes, the battery life could be a lot better. Yes, it could be faster. But it's still a device that even five years ago wouldn't have seemed possible. **–JACOB RIDLEY**

VERDICT

8

Ayaneo Air 1S

➕ **PULSE RACING** Specifications are incredible; Impossibly compact proportions.

➖ **TOTALLY FLATLINED** Performance is compromised; Battery life is awfully limited.

$1,029, www.ayaneo.com

SPECIFICATIONS

APU	AMD Ryzen 7 7840U
Cores/Threads	8/16
GPU	12CU AMD RDNA 3
Memory	32GB
Screen	5.5-inch AMOLED
Resolution	1920 x 1080
Refresh rate	60Hz
Battery	38Wh
I/O	2x USB 4.0 Type-C, 3.5mm audio, MicroSD
Dimensions	22.5 x 9 x 3 cm
Weight	405g

Even five years ago, the incredibly compact Ayaneo Air 1S would not have seemed possible.

The Prestige 13 Evo is a delightful little laptop, but comes at quite a price.

MSI Prestige 13 Evo

MSI aims for the premium ultrabook crown

MAKE NO MISTAKE: the MSI Prestige 13 Evo is a premium laptop. This almost feels like the blueprint for the perfect ultrabook, a supremely lightweight but still powerful laptop with a whole host of features and a sleek, minimalist aesthetic.

Coming in either white or gunmetal gray, the Prestige 13 Evo is equipped with almost everything you could want from an ultrabook, including a wide variety of physical ports, a large, responsive trackpad, and a selection of useful security features.

Weighing in at just 0.99kg (2.18lbs), this laptop sits in the same weight class as the popular LG Gram, and it's even lighter than the eminently portable M2 MacBook Air. With a 13.3-inch display and a thickness of just 1.7cm, it's phenomenally easy to pick up the Prestige 13 Evo and take it wherever you go. In fact, the AC adapter is also very lightweight, but you won't need to bring that everywhere, since this ultrabook also offers some impressive battery life.

My immediate comparison point for any compact ultrabook like this is the Dell XPS 13, which has long sat among the best laptops out there. MSI's laptop is actually a bit closer in price to the XPS 13 Plus, and it just about edges out Dell's competitor thanks to slightly better average performance and a more practical physical design.

While the stripped-down appearance might not appeal to everyone (the XPS 13 Plus has this one beat purely in terms of aesthetics), it's function over form here. The design is straightforward, with a robust screen hinge that gently angles the keyboard towards you when opened, and a large touchpad at the bottom. The exterior construction is plastic to further serve the goal of reducing the overall weight, but it doesn't feel cheap. Certainly, there's appeal in the straightforward design choices made by MSI.

As for performance, the 13th-gen Intel Core CPU at the heart of this laptop more than pulls its weight, giving you the option of some light gaming and creative work alongside the usual productivity tasks we test for on ultrabooks. I was impressed by the smooth, responsive user experience. General use is speedy and lag-free. You can open a dozen tabs in Google Chrome with Steam and Spotify running in the background, and not experience any slowdown whatsoever.

In synthetic benchmarks like GeekBench 6, the i7-1360P CPU demonstrates excellent single-core performance and solid multi-core results, putting it head and shoulders above its 12th-generation Intel counterparts. The SSD? It's merely adequate at about 1.65GB/s.

Naturally, the lack of a dedicated graphics card means you won't be doing any high-end gaming or 4K video editing tasks on this laptop, but that's fine—it's a small sacrifice to make for the incredibly thin and light design.

Thermal performance is also excellent; the Prestive 13 Evo has a large perforated section on the underside for venting excess heat, and the interior thermal solution clearly works well—it barely got warm in our testing process.

If you want an obvious negative, let's start with pricing. At $1,499 with apparently only one configuration, it's expensive, matching the XPS 13 Plus. At this price, the Prestige's relatively run-of-the-mill FHD+ display pales a bit in comparison to the 3.5K OLED screen of the Plus, not to mention paling in a literal sense, since it simply can't match the OLED's brightness and rich color density.

Other than the somewhat lackluster display and the mediocre SSD, there's a lot to like here. The lightweight design, strong Intel Core processor, and range of security features make it a highly practical device. Actually, I wouldn't mind hanging onto it forever. My own daily laptop is starting to look a little tired, and I'd happily ditch it for this classy compact computer. –**CHRISTIAN GUYTON**

VERDICT

8

MSI Prestige 13 Evo

➕ **PERFECTLY HONED** Impressively thin and light; Solid all-round perfomance.

➖ **QUICKLY DISOWNED** Disappointing screen; SSD is slow.

$1,499, www.msi.com

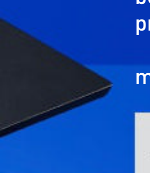

BENCHMARKS

	MSI Prestige 13 Evo A13M	**Lenovo ThinkPad X1 Carbon (Gen 10)**	**Acer Swift 14**	**LG Gram 16**
Geekbench 5.5 Performance (Score)	8,611	8,159	**11,375**	8,431
Video Transcode 4K to 1080p (Time)	10:41	12:33	**05:34**	17:00
HDD Transfer Rate (MBps)	744	1,462	1,102	**1,507**
3DMark Fire Strike (Score)	5,541	4,733	**5,659**	4,487
Civilization VI: Gathering Storm Benchmark (1080p, Medium) (fps)	31	**43**	28	19
Battery Life (hours)	12:03	08:48	11:43	**13:08**

Best scores are in bold. The Lenovo ThinkPad X1 Carbon and LG Gram 16 featured an Intel Core i7-1260P CPU, whereas the Acer Swift 14 boasted an Intel Core i7-13700H.

SPECIFICATIONS

CPU	Intel Core i7-1360P
GPU	Intel Iris Xe Graphics
RAM	16GB LPDDR5
Storage	1TB PCIe NVMe SSD
Screen	13.3-inch, 1920 x 1200, IPS
Ports	2x USB-C with Thunderbolt 4, 1x USB-A 3.2, 1x microSD, 1x HDMI 2.1, 1x 3.5mm combi audio jack
Wireless	Wi-Fi 6E, Bluetooth 5.3
Camera	1080p 30fps webcam, IR cam
Weight	0.99 kg (2.18lbs)
Dimensions	29.9 x 21.0 x 1.7 cm

© UNSPLASH

Asus Vivobook S 15 OLED

A lean, super-clean portable computing machine

A RELIABLE LAPTOP under $1,000 that doesn't sacrifice build, performance, or features? Tricky. But the Asus Vivobook S 15 OLED squeezes into that happy medium of offering efficient performance in a premium-looking chassis, all for well under $1,000.

Packing a 12th Gen Intel processor, a commendable 15.6-inch OLED display, and ample battery life, Asus did a fine job with this year's Vivobook line-up—or should I say, a clean job. Asus also brings its special "Antibacterial Guard" that claims to inhibit the growth of bacteria by more than 99 percent.

For the record, this is what Asus says about the feature: "The Asus Antibacterial Guard treatment consists of positively charged silver ions, and is durable enough for at least three years of normal usage. Household cleaning products and agents such as alcohol wipes and soap water won't easily wear away the silver-ion surface treatment, ensuring Asus Antibacterial Guard's long-lasting bacterial-inhibiting effect."

Despite it being a plastic build, that actually doesn't take away from the laptop's sturdiness or premium-looking Neutral Grey chassis. The Vivobook S 15 OLED also comes in Indie Black and Brave Green. Each has raised branding on the rear of the display outline and some racing-style stripes. That's unusual, as are the "#GoFurther" and "#BeFearless" slogans on the badges, even if that's unusual in a good way or in a cringeworthy way. You decide.

Open the lid, and you're treated to a wide Chiclet-style keyboard and a sizable trackpad. There are tiny details to spice up its design, including a bright orange escape key and the aforementioned racing stripe-type lines across enter. There are also spaced-out lines along the top of the touchpad, which is Harman Kardon's subtle touch on the laptop, along with a side-firing speaker system that's slightly better than the norm.

It's the display that steals the spotlight, however, with its "NanoEdge," slim-bezel panel and lay-flat 180-degree hinge. Oh, and 600 nits of peak brightness and 100 percent DCI-P3 coverage. It's a fabulous display, albeit not super-high resolution. Connectivity is good, too, with a decent selection of ports, including dual speedy Thunderbolt 4 ports. One of these is used to charge the laptop, so if you're keeping this laptop plugged in at home, you're still left with one port that supports external monitors, with up to 8K resolution and fast 40 Gbps transfer rates.

Size-wise, the Asus Vivobook 15 S OLED isn't the smallest or lightest laptop of its kind. The Samsung Galaxy 2 Pro 360 and Dell XPS 15 OLED are smaller in size, and the Vivobook leans towards the heftier side of the XPS 13 in terms of weight. But it's still a slim laptop that's easy to carry around in a laptop bag.

As for performance, the 12th Gen Intel Core i5-12500H CPU, 8GB of DDR4 RAM and 512GB SSD are a strong combo. Okay, sticking with DDR4 RAM instead of DDR5 compromises performance on paper. In practice, you'll never notice. That said, don't expect to go gaming; the Intel Iris Xe graphics are feeble. Even an entry-level last-gen GPU like the Nvidia GeForce RTX 3050 Ti will give the integrated Intel GPU a hammering.

Still, one upside of the low-power graphics is decent longevity from the battery. We extracted 7.5 hours of proper working, including web browsing, document editing, and content consumption. That's not far off genuine all-day battery life.

All of this means that the Asus Vivobook S 15 OLED excels as a capable work machine at an affordable price. There's good performance from the 12th Gen Intel chip, quality on-screen entertainment from the 15.6-inch OLED display, and you can (kind of) get through a full day of work. Admittedly, some elements, like the keyboard, trackpad, and webcam, betray this laptop's true value positioning. But for the money, it's a very nice machine.

–DARRAGH MURPHY

VERDICT

8

Asus Vivobook S 15 OLED

+ **OLE, OLE, OLED** Stunning 15.6-inch OLED display; Decent performance and battery life.

− **OH NO, OH NO, OH NO** Terrible webcam; Mediocre keyboard and trackpad.

$899, www.asus.com

SPECIFICATIONS

CPU	Intel Core i5-12500H
GPU	Intel Iris Xe
RAM	16GB
Storage	512GB SSD
Display	15.6-inch, 1920 x 1080, OLED
Ports	2x Thunderbolt 4, 1x UBA-A, HDMI 2.0, 3.5mm audio
Size	14.17 x 9.17 x 0.74 inches
Weight	3.97 pounds

© GETTY

The Asus Vivobook S 15 OLED delivers efficient performance and a premium-looking chassis for well under $1,000.

Asus ROG Zephyrus M16

Monster performance. Manic noise. Mega money

THIS VERSION OF ASUS'S ROG Zephyrus M16 isn't about affordability; it's all about maximum performance in a modest chassis. Yes, we're looking at the daddy, the RTX 4090 version, though not to the extreme 175W level that the likes of Razer have allowed this GPU to achieve in their systems.

Instead, you get the slightly more parsimonious 150W RTX 4090 spec, which is still technically at the top of its potential TGP rating, but without the extra 25W leeway manufacturers are allowed to add on top. Of course, this isn't the same as the desktop version of the RTX 4090. It's actually using a completely different GPU, the AD103 chip as seen in the desktop RTX 4080, right down to the core, cache, and memory configuration.

Is it still a high-performance GPU? For sure. But ultimately, it's an RTX 4090 in name only. Anyway, the Zephyrus M16 also comes sporting one of Intel's latest Raptor Lake CPUs. The Core i9 13900H is a 14-core chip with 20-threads. In Intel's hybrid parlance, that means six Performance-cores with Hyperthreading, and eight Efficient-cores without. Put simply, this CPU is unlikely to ever bottleneck the monster Nvidia GPU in games.

Its 5.4GHz claimed boost clock speed isn't some token figure for the marketing material, either. You will actually see that level of CPU performance in gaming—well, you will if you let the processor run that fast. As with a lot of gaming laptops from this generation of machines, if you want to access the full performance potential, be prepared to let it run very hot and incredibly loud.

The cooling inside the Zephyrus M16 is actually good enough that both the RTX 4090 and the Core i9 CPU are able to run at top levels. The trade-off is unreasonable noise levels. Maybe, just maybe, a really good set of noise-cancelling headphones would make the awful din this laptop makes on full reheat tolerable. But the cacophony really is remarkable.

At least it's fast when it's making all that noise. Despite trading that extra 25W on the RTX 4090 TGP with the Razer Blade 16, the Zephyrus M16 is still able to deliver frame rates almost on a par. In some CPU-bound games at 1080p, you can even see it pulling ahead, even though it doesn't have as powerful a processor inside it.

Arguably, the slightly less manic "Performance" setting in the Armory Crate software gives a better balance of fan noise and gaming frame rates over the full "Turbo" madness. Whatever, the configurable buttons above the keyboard can be set to cycle between power settings. Being able to easily switch between Performance, Turbo, or even "Silent" makes for a positive everyday PC-using experience, and allows you to quickly decide what you noise tolerance levels are like at any given moment.

As for the Zephyrus's screen, it's a great middle ground between the demands of the 4K panel in the Blade 16 and an older high-refresh rate 1080p display. The 2,560 by 1,600 native res isn't nearly as demanding on the GPU as 4K, but at 16 inches, you still get a really tight pixel pitch. The mini-LED backlight makes for fabulous HDR performance, too. Actually, it's probably the best thing about this laptop. It really is that good.

We also prefer the overall styling here to that of Razer Blade 16, despite the Zephyrus's slightly chonkier chassis. It's still impressively slight, given the performance, and the lifting action of the screen hinge remains a smart way of increasing airflow. All told, it's a bit more angular and more traditionally 'gamery'.

In summary, we'd probably pick the Asus rig over Razer—at least we would before you factor in price. Yes, this is the top spec for the Zephyrus M16, but $4,000 for a gaming laptop, even one this good, is borderline unconscionable. **–DAVE JAMES**

VERDICT

8

Asus ROG Zephyrus M16

➕ BEAUTIFUL BEAST Mini-LED screen rocks; Great gaming performance.

➖ HORROR SHOW Extremely loud; Incredibly expensive.

$3,999, ww.asus.com

BENCHMARKS

	Asus ROG Zephyrus M16	Razer Blade 16
Cyberpunk 2077 1080p RT (fps)	60	**64**
Cyberpunk 2077 1440p RT (fps)	39	**42**
F1 2022 1080p RT (fps)	101	**106**
F1 2022 1440p RT (fps)	73	**82**
Metro Exodus Enhanced 1080p RT (fps)	**113**	99
Metro Exodus Enhanced 1440p RT (fps)	**87**	86
Hitman 3 1080p (fps)	240	**270**
Hitman 3 1440p (fps)	188	**207**
Cinebench R23 9pts)	19,689	**28,476**

Best scores are in bold. Comparison system in the Razer Blade 16 with 175W Nvidia RTX 4090 graphics, Intel Core i9 13950HX CPU and 32GB RAM.

SPECIFICATIONS

CPU	Intel Core i9 13900H, 14 cores, 20 threads
Boost clock	5.4GHz
GPU	Nvidia RTX 4090 (150W)
VRAM	16GB
System Memory	32GB DDR5-4800
Storage	2TB NVMe SSD
Screen	16-inch, 2560 x 1600, 240Hz
Ports	1x Thunderbolt 4, 1x USB-C, 2x USB-A, 1x HDMI 2.1, 3.5mm audio
Dimensions	13.98" x 9.69" x 0.90"
Weight	5.07 lbs

© UNSPLASH

Asus's latest Zephyrus sports a glorious mini-LED screen, powerful CPU, and of course, a beefy GPU.

SUBSCRIBE TODAY

Find your next great read

On iOS & Android!
Mac|Life
iOS 17
GET YOUR UPGRADE!
MAXIMUMPC
BEST VALUE GAMING PC
PC GAMER
THE TOP 100 PC GAMES 2023
STARFIELD

Subscribe today and get instant access on your iPad, iPhone or Android device.

www.magazinesdirect.com/B3XMPC

Logitech Pro X Superlight 2

Hold it down or it'll float away

EVERYONE KNOWS that gamers want two things: migraine-inducing RGB, and over 9,000 buttons per mouse. Or do they? Logitech begs to differ, and its Pro X Superlight 2 aims squarely at one all-important metric: minimal weight.

At just 60g, it's 3g lighter than its predecessor. Impressively, it clocks up this lighter weight without cutting any holes into its frame. In fact, it's 8g lighter than the holepunched SteelSeries Aerox 3 Wireless Ghost, and manages to do so without looking like an RGB cheese grater (let's not give Razer any ideas).

Instead, you'll find a comfortable, sturdy frame that doesn't mess with the classics. There's no outlandish swoops and curves, no overabundance of buttons, and certainly no RGB. The pared-back design keeps things simple, and doesn't blind you with flashing lights. There is a pink color option, sure, but it doesn't glow in the dark.

OK ZOOMER

The reduced weight is not the only upgrade over the previous version. There's a new Hero 2 sensor, the maximum DPI has increased from 25,600 to a frankly obscene 32,000, while the polling rate has doubled to 2,000Hz. Gone is the prehistoric microUSB port (replaced by USB-C), battery life is up from 70 hours to 95, and the once-mechanical switches are now a hybrid of optical and mechanical parts. There are many small changes, but they add up.

On the bottom are two PTFE feet that help the Superlight 2 glide effortlessly over your mouse mat. Combined with the incredibly low weight, it means you can zip around the place faster than *Baldur's Gate 3*'s Astarion bedding... well, just about anything. It's not quite as smooth as the Razer Cobra Pro, but it's not far off.

The effect on gameplay is superb, and it excels in fast-paced shooters, where the new Hero 2 sensor is quick, accurate, and wonderfully consistent. In slower games where you don't necessarily need the low weight and gliding smoothness, you still appreciate them over heavier, clunkier alternatives.

The hybrid optical-mechanical left- and right-click buttons are firm and solid. They're also loud, and I could easily hear them over music while wearing a pair of over-ear headphones. This is one mouse that really wants you to know that it's doing a great job of clicking, bless its heart.

On the side are two customizable buttons, which are quiet and soft to press. The scroll wheel, meanwhile, is satisfyingly lightweight—I could just scroll it up and down all day (but instead I have to "work" and "play games"). Clicking it, though, feels dull and heavy, like trudging through thick mud. At least it's quiet.

BATTERY FOR DAYS (AND WEEKS)

One benefit of the lack of RGB is the extensive battery life you get with the Superlight 2. Logitech quotes 95 hours, and while that's not the longest innings we've ever seen, it should last you a couple of weeks before needing to be juiced up again. After just over 35 hours of solid game and work usage, it was still going strong at 89 percent battery in my testing, so I'd expect to get much more than the quoted 95 hours.

To check the battery life, you've got to open up Logitech's G Hub app, which isn't great. Assigning macros is confusing, and while there is a ton of depth and control when it comes to DPI options, my app had a habit of resetting my DPI whenever I opened a full-screen game.

Speaking of DPI, the Superlight 2 lacks any way to change this on the mouse itself, so you have to wade into G Hub every time you want to adjust it. It's all part of the understated physical design, which I suppose could also be described as a touch unimaginative.

But I'll take unimaginative every day of the week if it performs this well. The Superlight 2 does the basics, and it does them really, really well. Its cloudlike lightness makes it ideal for fast, intense games, but it's a brilliant pointer for much more than just that. With super-smooth gliding and an excellent sensor, it'll be a fantastic choice for a whole heap of gamers. **–ALEX BLAKE**

VERDICT

9

Logitech Pro X Superlight 2

⊞ **GREEN LIGHT** Seriously lightweight; Slides nicely on a mouse mat; Phenomenal battery life; Pleasingly smooth scroll wheel; Sturdy construction.

⊟ **RED LIGHT** Middle mouse click feels dull and heavy; G Hub software isn't great; No DPI button; Design isn't particularly exciting.

$159, www.logitech.com

SPECIFICATIONS

Connection	2.4GHz wireless (Logitech Lightspeed adapter), USB-C to USB-A wired
Sensor	Hero 2
Acceleration	40G
Speed	500 IPS
Battery	95 hours
Polling rate	2,000 Hz
Customizable buttons	3 on top, 2 on the left-hand side
Weight	60g

Phanteks NV7

What happened to all the wires?

PCS KEEP GETTING more powerful and indeed more power hungry. That means bigger components like the monster Nvidia RTX 4090 GPU. If you're spending $1,500 on a graphics card, there's a decent chance you might want to show off your hardware.

Enter the Phanteks NV7. It's a full tower that's slathered with glass, mesh, and RGB. If that hardly seems revolutionary, the NV7 manages to look and feel fresh, with lots of customization, and a $220 price that's reasonable with the proviso that you have to bring your own fans.

Still, the NV7 has the most cooling space of just about any case we've reviewed. This is a big boy, measuring 20.9 by 10 by 23.1 inches, making it roughly two inches taller than the already beefy Lian Li O11 Dynamic XL. For the record, it can house GPUs up to 17.7 inches and CPU coolers up to 7.28 inches.

12 fans can fit in the NV7, and that's without sacrificing any storage space. Phanteks sent along a few sets of its new D30-120 fans, which don't come as standard, but sport an internal RGB ring and clip together like the Lian Li's Unifans, reducing the need for fiddly cables and connectors.

Regardless, installing fans is a slick process. The side fan mount is held in by one thumb screw at the top of the case, while the bottom tray can be rotated a few degrees to angle the bottom fans forward if you want. There's also a GPU anti-sag bracket included that's relevant to thermals, but not necessarily in a good way. While it is adjustable to fit the width of most GPUs, it wobbles like crazy and restricts airflow.

Like most of Phanteks' cases, the RGB is sprinkled here and there, and never overwhelming. The case features an RGB strip that runs along the edge of the motherboard tray. It's all very customizable, thanks to the four control buttons at the top of the chassis.

That said, one area where the design feels a bit lazy is on the rear side panel, which needs to be screwed in from the top of the case with a tiny screw that would be all too easy to lose.

After removing the rear side panel, you'll see the unique storage caddies in the NV7. The design is clean and modular, and very capacious, with support for six 2.5-inch and one 3.5-inch drive, or two 2.5-inch and two 3.5-inch drives.

By now, you're probably asking where the power cable connects. Everything from the power cable to the motherboard IO is hidden behind a mesh door, and the cables are looped out from the bottom, adding an extra step to the build.

As for IO, it's a little less abundant than usual for a full tower. You just get two USB 3.0 Type-A ports and one USB 3.2 Gen 2 Type-C port, alongside a headphone/mic combo jack set at an angle on the bottom front. The power and RGB buttons live up top, and while the controls work well enough it would have been nice to have some ports up there as well.

In testing, the NV7 didn't suppress fan noise terribly well, but the upside is excellent thermal performance. However, if pure air cooling isn't sufficient, the NV7 also pairs well with Phanteks' many custom water-cooling components, with a fill port up top and a drain port in the bottom rear, plus loads of space for multiple radiators and tubes for all major components.

The $220 price is appealing, too. Okay, you'll have to spend more to bring your own fans. But chances are if you can afford the kind of high-end components and custom cooling options this case was designed for, you can afford to pay a little extra for some fans. **–MYLES GOLDMAN**

VERDICT **8**

Phanteks NV7

+ TOWERING ACHIEVEMENT Novel cable-routing design; Tons of RGB customization.

− TALL ORDER No fans included; Mediocre noise suppression.

$220, www.phanteks.com

SPECIFICATIONS

Type	ATX Full-Tower
Dimensions	20.9 x 10 x 23.1 inches
Max GPU Length	17.7 inches
CPU Cooler Height	7.3 inches
Internal Bays	6x 2.5 and 1x 3.5-inch
Expansion Slots	8
Front I/O	1x USB 3.2 Gen 2 Type-C, 2x USB 3.0
Fans	2x rear 120mm, 3x top 120mm, 4x side 120mm, 3 x bottom 140mm
Weight	37 pounds
Warranty	5 years

Corsair K70 Max

Is this the fastest keyboard ever?

AT LAST, Corsair has brought Hall-effect magnetic key switches to its famed K70 keyboard design. Every switch in the K70 Max is fitted with one of Corsair's new MGX switches, hoping to tap into that magnetic-mechanical magic that has become a bit of a trend lately in the world of enthusiast gaming keyboards.

Magnetic key switches have a few major benefits over standard mechanical switches. Chief among them are adjustable actuation points, making for a customizable typing experience that caters to both gamers and typists, depending on how you set them up.

The Corsair's MGX switches offer actuation between 0.4mm and 3.6mm. That makes them snappier than MX Speed switches on the one hand, which actuate at 1.2mm, and seriously tough to mistype on the other.

For most gamers, it's worth trying the lowest actuation. Only the lightest touch will actuate the switch, and while that's not great if you're a clumsy typer, it means your actions are relayed in-game post-haste. Anyway, paired with the 8,000Hz polling rate and adjustable reset points, the K70 Max is incredibly rapid.

Another intriguing feature is dual-point actuation. This takes getting used to, but the idea is assigning two actions to a single key. The actuation points for these dual-action keys can be set in the Corsair iCUE app, but ideally you'll have one action actuate half-way down the switch, the other when it bottoms out.

As for the typing experience on the K70, the board feels decent. It includes two sound-dampening layers, but compared to the Asus ROG Strix Scope II 96 Wireless, it's not only the louder of the two, but not as pleasant to type on. The rattle of the keys is pretty noticeable versus the Scope II's lubed switches.

Still, the K70 Max is a good fit for competitive gamers, not just for the speed, but the toggleable Tournament Mode. This locks the lighting to a single color, removes all custom actions, flips the keyboard to a basic profile, and delivers rapid access to its purest mode.

There's also no faulting the Corsair for looks. The per-key RGB is about as good as you'd expect from a company known for slapping LEDs on everything. But it's actually the board's ability to nail white light that brings out the colors. All the RGB lighting is customizable via the iCUE app. It's one of the better applications for putting all those options in one place but still keeping thins easy to navigate.

The build quality is impressive, too, especially on the included wrist rest. But be warned, this is a chunky keyboard. Corsair loves a thick board, and with the added height of the top row, fitted with a volume wheel and some shortcut keys, it's a desk hog. The tiny the num and caps lock indicators next to the Corsair logo in the center of the board seem lost, too.

Then there's the price. For what is a solid iteration on a classic gaming keyboard, the K70 Max at $230 is expensive. That's a lot next to enthusiast-mimicking boards such as the ROG Strix Scope II 96, a cheaper option with a superior general typing experience.

If Corsair could nail down that sort of robust typing feel and satisfying clack alongside the especially quick competitive gaming features, the K70 Max would be the board to beat for both gamers and typists. For now, the K70 Max is a high-end keyboard that will appeal to those chasing the snappiest response. But it's too expensive, and with the noise and slight lack of tactility, a little too flawed for an unqualified recommendation. **–JACOB RIDLEY**

VERDICT
8

Corsair K70 Max

+ SPEED DEMON Adjustable magnetic switches; Uber-fast 8,000Hz polling rate.

− OFF THE PACE Pretty price; A little noisy; Only okay for typing.

$230, www.corsair.com

SPECIFICATIONS

Switches	Corsair MGX (linear)
Connectivity	Wired (USB Type-A to Type-C)
Backlighting	Per-key RGB
Onboard memory	8MB
Polling rate	Up to 8,000Hz
Media controls	Yes, volume wheel and shortcut keys
Dimensions	17.4 x 6.54 x 1.54 inches
Weight	3.06 pounds

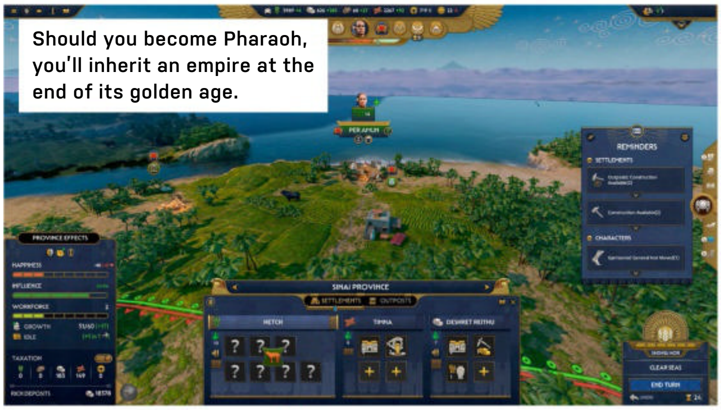

Should you become Pharaoh, you'll inherit an empire at the end of its golden age.

Zoom in on the action, and you hear maces thudding into the skulls of their victims.

TURN BASED STRATEGY

Pharaoh's central belligerents have none of the personality of *Three Kingdoms*' colorful generals.

Total War: Pharaoh

Back to basics, but the most complex entry yet

CREATIVE ASSEMBLY'S chosen setting for *Total War: Pharaoh* sits right upon the precipice of military history. The Battle of Kadesh, which took place about 100 years before *Total War: Pharaoh* begins, is the earliest pitched battle for which we have records of tactics and formations.

Playing as an Egyptian, Canaanite, or Hittite general, your goal is to conquer enough of the land around the eastern Mediterranean that your name is forever carved into the stele of history. In this, becoming Pharoah is merely the start of your aspirations. Around turn 12 of the campaign, you unlock a system known as "Power of the Crown". This lets you choose to follow the path of becoming either the Egyptian Pharaoh or the Great King of Hatti. These positions come with slightly different benefits, but they're attained in the same way—winning a civil war.

The impact this system has is significant. For starters, it makes *Total War*'s early game more immediately engaging, as you can have a crack at becoming Pharaoh as soon as turn 12. It also means that progression isn't necessarily a gradual, one-way accrual of power. You might become Pharaoh early on, but then you'll need to maintain that position by fighting off pretenders to your throne.

This system is nothing like as involved as Crusader Kings' elaborate political simulation, playing out more like a minigame that refreshes every in-game year. But it is fun to play around with, and can tie into the broader game in some interesting ways.

Should you succeed in becoming Pharaoh, your problems will be far from over, as you'll inherit an empire at the end of its golden age. There are five resources: food, wood, stone, bronze, and gold. Food and bronze are crucial for maintaining armies, wood, and stone for constructing buildings, while gold plays a role in pretty much everything.

Simply getting hold of all these resources can be tricky. Gold and bronze deposits are scarce on the map, while stone has a hard limit on how much you can mine across the entire game. This makes trading and bartering more vital than in any other *Total War*. Simply maintaining a decent-sized army requires you to solve a whole bunch of logistical problems.

So, Creative Assembly has built a sumptuous, diverse and thematic campaign. There's a whole bunch of other stuff we haven't mentioned, like how praying to different gods can provide different bonuses, and the victory condition-like Ancient Legacies that can spur you on your way to victory. But if there's a weakness, well, it's the battles.

To be clear, Pharaoh's battles are not bad. They are classic *Total War* affairs that combine robust tactical foundations with brutal physicality. Walls of infantry crashing into one another, arrow hailstorms thudding into the sand, chariot charges sending footsoldiers flying—all that good stuff. *Pharaoh* also does what it can to add some new ideas. Weather conditions are more diverse and dynamic. Being largely a desert country, Egypt's weather is less rain and fog, and more sweltering heat and choking dust storms.

Ultimately though, there is only so much Creative Assembly can do to spice up the ranks of such an ancient

© SEGA

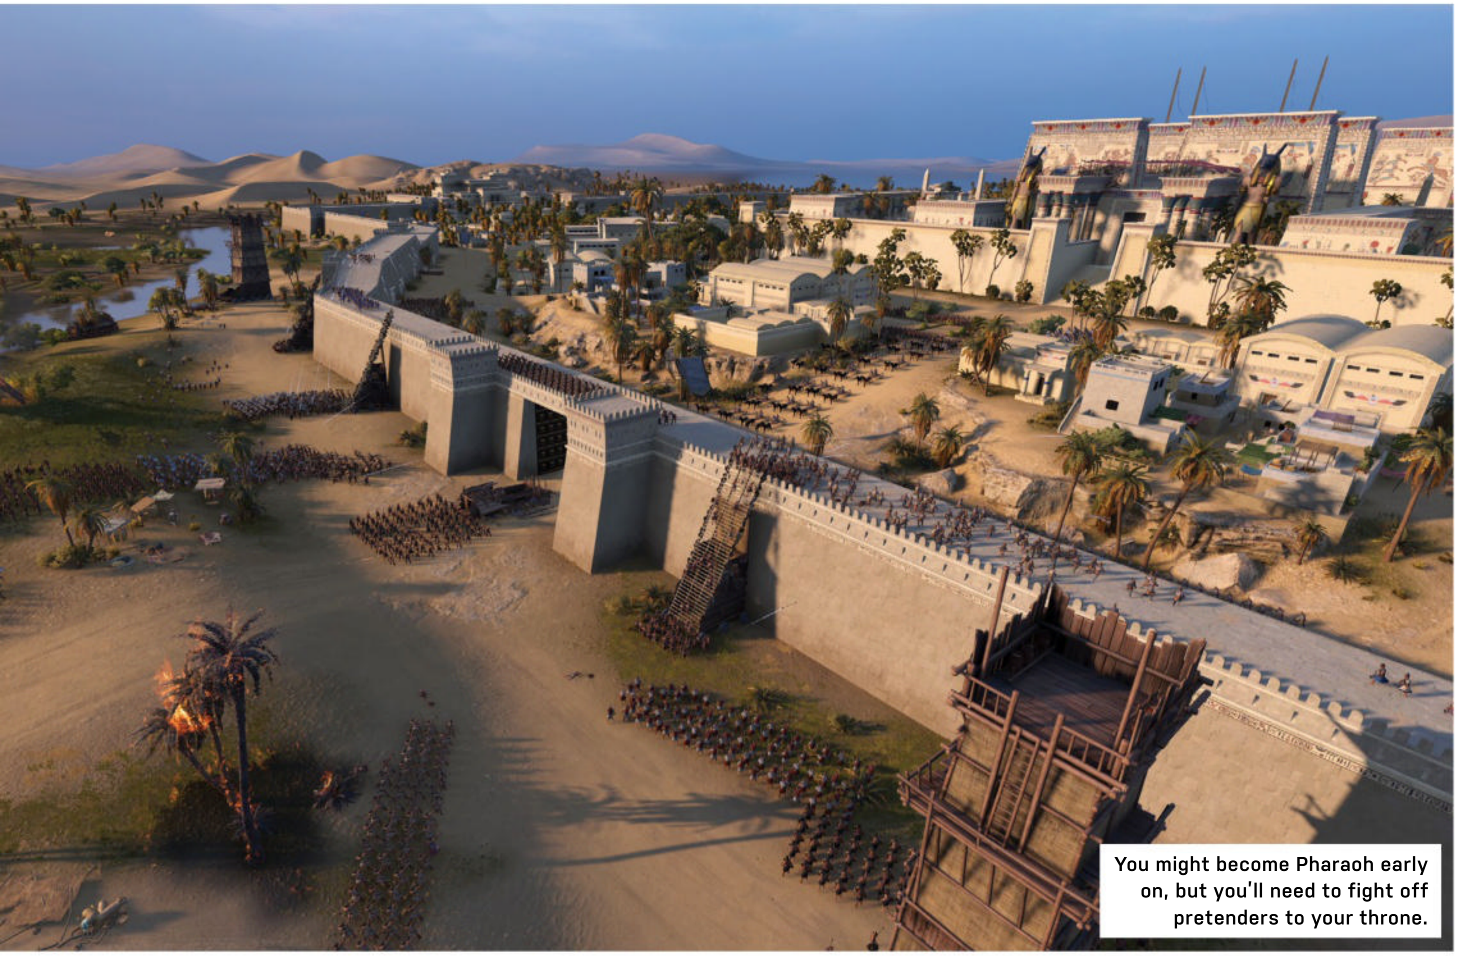

You might become Pharaoh early on, but you'll need to fight off pretenders to your throne.

civilization. There are many faction-specific types of swordsmen, axemen, spearmen, chariots, etc. But units are all slight variations on these categories. After controlling the wildly imaginative factions of *Warhammer*, and the pseudo-magical armies of *Three Kingdoms*, going back to slings and arrows is a bit deflating.

Still, *Pharaoh* is a success. This is a much harder setting to make work in *Total War*'s context than more familiar locales like Rome or feudal Japan. The battles may be simplistic, but they have that *Total War* magic, and Creative Assembly has built an evocative and exciting campaign around them. –RICK LANE

VERDICT
8

Total War: Pharaoh

RAMESSES THE GREAT Wonderfully rich and complex campaign; Immediately engaging.

XAKHENATEN THE AWFUL Battles are a bit basic; AI commanders lack character.

RECOMMENDED SPECS CPU, Intel i5-6600 / Ryzen 5 2600X. GPU, Nvidia GeForce GTX 1660 Ti / AMD RX 480. RAM, 8GB.

$60, pharaoh.totalwar.com, rated 16

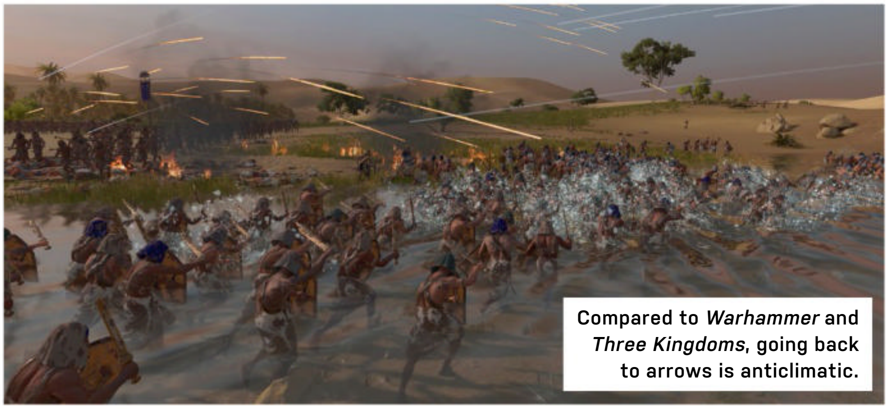

Compared to *Warhammer* and *Three Kingdoms*, going back to arrows is anticlimatic.

The battle scenes successfully retain that *Total War* charm.

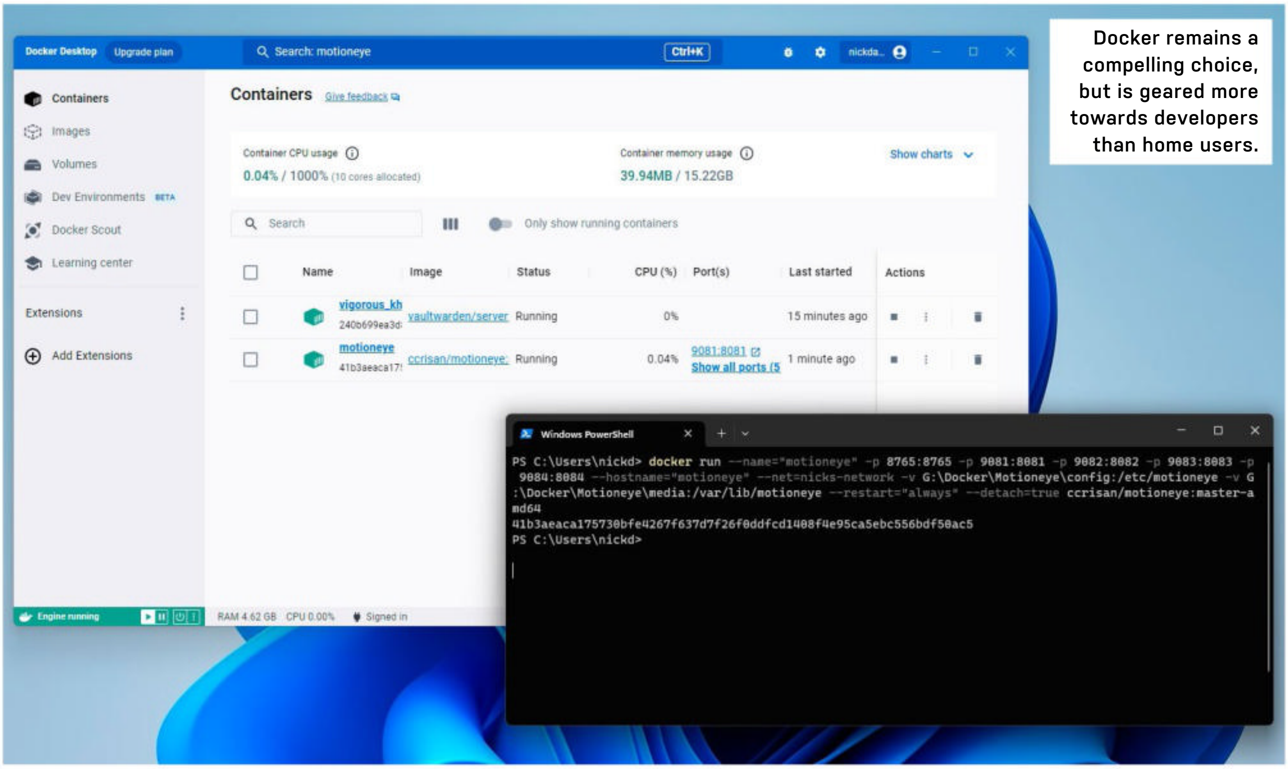

Container CPU usage 0.04% / 1000% (10 cores allocated)
Container memory usage 39.94MB / 15.22GB

☐	Name	Image	Status	CPU (%)	Port(s)	Last started	Actions
☐	vigorous_kh 240b699ea3d:	vaultwarden/server	Running	0%		15 minutes ago	
☐	motioneye 41b3aeaca17!	ccrisan/motioneye:	Running	0.04%	9081:8081 Show all ports (5	1 minute ago	

```
PS C:\Users\nickd> docker run --name="motioneye" -p 8765:8765 -p 9081:8081 -p 9082:8082 -p 9083:8083 -p
 9084:8084 --hostname="motioneye" --net=nicks-network -v G:\Docker\Motioneye\config:/etc/motioneye -v G
:\Docker\Motioneye\media:/var/lib/motioneye --restart="always" --detach=true ccrisan/motioneye:master-a
md64
41b3aeaca175730bfe4267f637d7f26f0ddfcd1408f4e95ca5ebc556bdf50ac5
PS C:\Users\nickd>
```

Docker remains a compelling choice, but is geared more towards developers than home users.

Docker vs Podman

What's best for running virtual machines for free?

NOT EVERYONE NEEDS a full-blown virtual machine to run certain types of software—Docker is a lightweight alternative that uses containers to wrap up Linux applications and dependencies in a minimal environment. Readers of this magazine will have seen our many tutorials for Docker over the past few years and got a good idea of how it works.

But Docker isn't perfect—it was originally designed for development rather than production-ready use, which has led to question marks over security and stability. Luckily, in 2015, Docker created the Open Containers Initiative (OCI), which has seen several alternatives emerge that can run the same containers. These include Podman, aimed more at production environments (and perfect for home users).

Developed by Red Hat, it strives to remain broadly compatible with Docker while employing a different means of managing containers to improve security and stability. Podman shares much of Docker's command-line syntax (simply prefix commands with 'podman' rather than 'docker'), making it relatively easy to switch from one to the other—at least for simple containers.

So how does Podman work differently? Docker runs a separate process (known as a 'daemon') to manage all running containers, which requires root access and, if it goes down, takes all running containers with it. Podman is 'daemonless', with each container allocated its own self-contained process. Better still, all containers are set up to be 'rootless' by default, so they don't have access to sensitive parts of the system.

Docker can be configured to run rootless, but it's a convoluted process. It's also true that some containers won't run in a rootless state, but Podman supports a simple 'set –rootful' flag for those containers that need elevated access.

INSTALLATION & DEPLOYMENT

You can install both Podman and Docker on Windows, macOS, and Linux. Those looking to deploy Docker in Windows or on the Mac must install the Docker Desktop version, which includes a GUI. One immediate disadvantage for Podman wannabes is that there's no official repository for Ubuntu, forcing users to make do with outdated builds from the distro's own repos or chance their arm with unofficial repos.

There are no such problems on Windows: in both cases it's a straightforward affair, with the Linux WSL2 subsystem installed if required during setup. From here, Docker is ready to go, while Podman requires a couple of extra commands ('podman machine init', followed by 'podman machine start').

As we've said, the pair share a similar command structure—as the interloper, Podman is clever enough to ape most of Docker's commands, and has been engineered so that Docker API tools and clients can use it as if it was Docker.

All core functionality is the same, despite Podman's daemonless approach. That said, this difference has created issues in the past, so it's important to have an up-to-date build thanks to Podman's rapid development over the past few years. For example, Podman 4.0 introduced support for creating custom networks in rootless mode,

© DOCKER

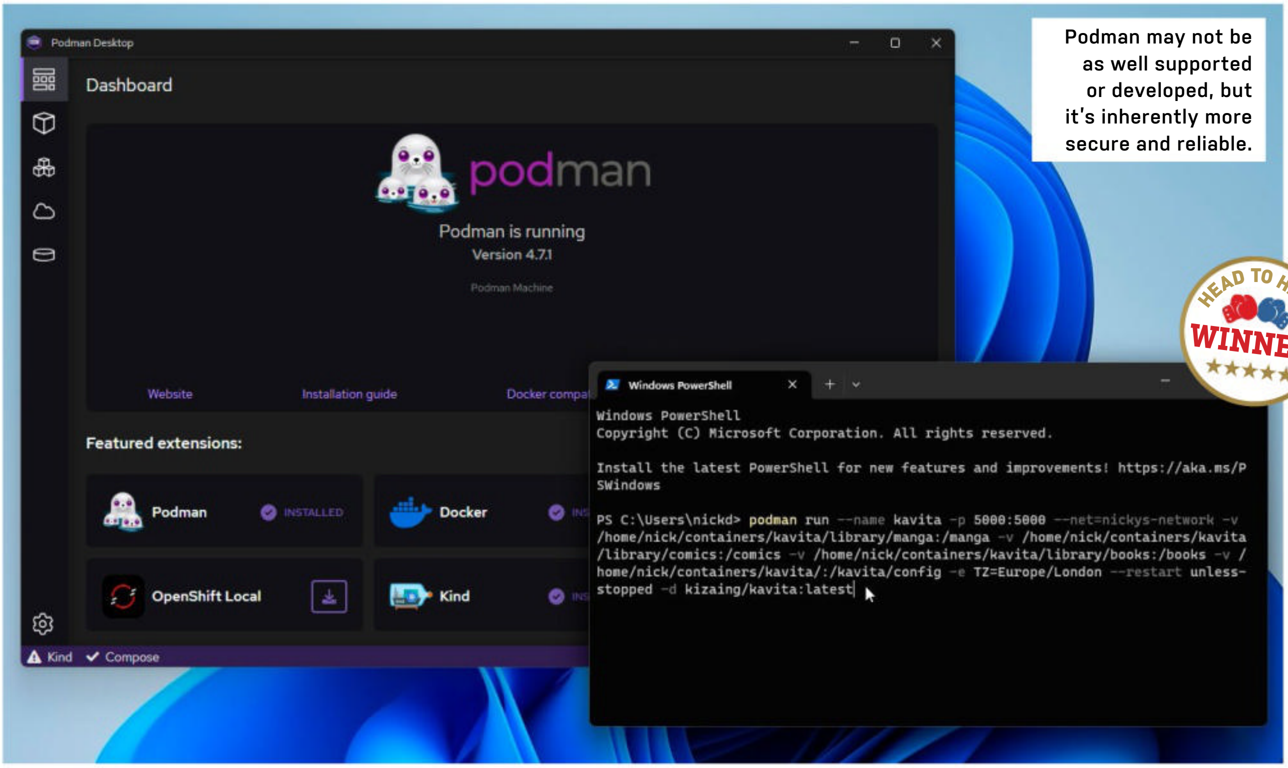

```
Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

Install the latest PowerShell for new features and improvements! https://aka.ms/P
SWindows

PS C:\Users\nickd> podman run --name kavita -p 5000:5000 --net=nickys-network -v
/home/nick/containers/kavita/library/manga:/manga -v /home/nick/containers/kavita
/library/comics:/comics -v /home/nick/containers/kavita/library/books:/books -v /
home/nick/containers/kavita/:/kavita/config -e TZ=Europe/London --restart unless-
stopped -d kizaing/kavita:latest
```

Podman may not be as well supported or developed, but it's inherently more secure and reliable.

© PODMAN

which is essential when trying to connect containers to the outside world through a reverse proxy.

In keeping with the Open Containers Initiative—and highlighting how Docker and Podman broadly do the same thing—both tools also allow you to package up containers so they can be migrated from one application to the other.

When it comes to divergence, Docker offers Swarms, which are described as a "container orchestration platform". It allows multiple hosts or servers running Docker to work in tandem with each other. That's overkill for most basic servers, and while Podman doesn't support Swarms natively, it supports Kubernetes, which offers a similar feature.

Podman does, however, support Docker Compose, which is designed to allow users to run applications within multiple containers, such as Nextcloud AIO, as featured in the October 2023 issue. However, Podman also provides an alternative in the form of Pods, which provide a shared space for related containers to reside. The downside is that the feature isn't extensively documented, which highlights one of Podman's weaknesses: relative lack of support compared to the long-established Docker.

Managing both Podman and Docker from the command line is similar, but for those who prefer user-friendly interfaces to the command line, both desktop tools are capable affairs, if lacking a few key features. Docker Desktop has the edge here—first, its universal search tool allows you to search online repositories for container images using keywords; Podman Desktop relies on you knowing the exact name and build you want. Second, Docker Desktop exploits one advantage of managing all containers within a single daemon by allowing you to allocate a specific amount of system resources to share across your containers.

Both tools also support other web-based frontends—Portainer will run on both, for example, although it's geared more towards Docker, while Podman users can install a plugin to view and manage containers through the Cockpit web UI in Ubuntu Server.

Which one should you choose? Docker is better established, and if you're currently using it, you may find it too much hassle to switch, although once Podman is up and running, migrating containers across is straightforward. Docker's main appeal may be with those who like to create and modify containers from scratch, but its extensive support still makes it a compelling choice for newcomers.

Podman has come a long way since its initial release, when its lack of features made it a hard sell. Now, the only thing counting against it is the relatively sparse documentation and support. However, if security and stability (along with reduced overheads) are your thing, Podman's daemonless approach makes it more compelling in our eyes. Put it this way, as soon as we can source a reliable up-to-date Ubuntu repo, we'll be switching our server over. Stay tuned. **–NICK PEERS**

VERDICT

8

Docker

ON THE PODIUM Long-established; Massively supported; Widely compatible.

IN THE DOCK Default setup requires root access; Not optimized for home use.

RECOMMENDED SPECS 64-bit CPU, 4GB RAM, WSL2.

Free and open source

VERDICT

9

Podman 4.7.1

POD-U-LIKE No daemon means better security, performance and stability; Rapid development.

POD PERSON Documentation and support not as extensive as Docker; No up-to-date Ubuntu build.

RECOMMENDED SPECS 2 physical CPU cores, 6GB RAM, WSL2.

Free and open source

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> Print natives
> Laptops or claptrap?
> OLED, where art thou?

No love for digital

I've been a subscriber to the magazine for over 20 years, and would read almost every issue of *Maximum PC* cover to cover within a week of receiving it. Since you went to a digital format, I struggled to get through the first issue. It took me almost two months to find the time to finish it. Since then, I haven't read another issue.

I realize that I'm a dinosaur, as I had my computer science degree before most of your readers were born, but I still build at least one PC per year, and like to keep my knowledge fresh, so I consider my expertise to be far above that of the average PC user.

I did enjoy reading your recent history of computing articles, and laughed when I realized that I either owned, or at least programmed on, most of those PCs.

I thought of printing the entire magazine to read it that way, only to discover that I'm only allowed to print a few pages at a time.

It pains me to say it, but I can't justify spending money on a product that I no longer use. It's not the content, it's the format. If you can at least make it so we can print the entire magazine, I would change my mind.

I do like that we now have access to back issues, but I still have them on my bookshelf, so it's not that much of a plus for me, personally. **– B. Frost**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: First of all, thanks for your letter detailing your frustrations with the digital version. I'm sure you're not the only person out there who prefers reading a physical product to something on a screen. There are times when I feel like all of my time looking at electronic displays is going to catch up with me, and one day I'm going to wake up with the need for really strong prescription glasses, or a problem with migraines, or worse.

As I'm sure you can appreciate, the pandemic through 2020-22 accelerated the decline of print, while distribution costs also went up, and the cost of living meant that subscriptions to luxuries like magazines were one of the first things to go in peoples' budgets. Thankfully, our publisher made the difficult decision to cut costs and move to digital this year, so that the brand and the people who work on it were protected, which is something that I'm immensely grateful for. We're now thriving again as a digital publication, and able to produce the content that you know and love as we have for decades now.

It sounds as though I might come from a slightly different generation to yourself (I'm 41), but I'm someone who has absolutely loved magazines since I was a teenager. That said, the idea of picking a physical copy of a title I love is something I only do on the rare times that I take a beach holiday, or am waiting for my dentist appointment. For me, there's nothing more relaxing and luxurious than picking up a magazine and reading it on the sofa or in bed, but I now do so on a digital subscription app like Readly on my Samsung Galaxy Tab S9 (covered in the November 2023 issue on page 75). That way, I can switch between my favorite magazines like *EDGE*, *PC Gamer*, and *Play* magazines with ease, and just read as much as I like of each.

Your point about being able to print out pages is valid though, and something other readers have asked about—I will bring this up in my next meeting with the management team here at Future.

Do laptops still suck?

I was reading through some back issues of *Maximum PC*, and in an issue from early last year, I saw that Jeremy Laird wrote an article saying he thought that gaming laptops suck. I then saw that new editor Guy Cocker issued a rebuttal against this, saying that he'd spent a lot of time during COVID using his HP gaming laptop to work during the day and play in the evenings.

Personally, I agree that gaming laptops do suck for pure gaming and

↘ submit your questions to: editor@maximumpc.com content creation, at least when they're compared to a *Maximum PC*-inspired desktop. But it's definitely a different market and usage model. I can't carry my desktop and set it up in my airline seat. It might also be a bit awkward setting it up in a coffee shop, and I don't have to pack a trunk-sized pelican case to take it to the hotel.

Simple LAN parties are also easier when using a laptop, logistically speaking. I have to travel a lot for work, and have been on the quest for the perfect laptop for around 25 years now. My first laptop was a Sony Viao with floppy disks and swappable batteries. I have one word to sum that particular quest up: compromise. I have used everything from a 17-inch desktop replacement (Ouch! My shoulder!) to a 13-inch road warrior (Please load this game before I have to leave...).

You always end up having to give up something and make trade-offs in order to get hold of what works for you. When it comes to pure power, you give up portability, and vice versa. Thin and light? Well, you probably have soldered memory, and maybe even the SSD is too. Forget about adding a second data drive.

My current travel buddy is an Asus ROG Zephyrus G14 with the Radeon RX 6800 S, and an 8TB SSD that I slapped in myself. Since my game library is so big, I'm pretty sure it has developed its own consciousness. I was a little leery of a Radeon graphics card, since they haven't had that much of a presence outside of Apple products on the laptop front, but it turns out that I've been pleasantly surprised with it.

I also love the MUX switch feature. I use the integrated graphics, and can power it on an airplane with the in-seat plug and a 100W PD charger while I play simple strategy games. It still runs hot though, which is my first complaint. The other is Win11, but I guess that I will have to learn it sooner or later. Anyway, thanks again for a great magazine, and thanks for listening. **– C. Von**

Samsung's Galaxy Book3 Ultra is one sleek gaming machine.

© SAMSUNG

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: Ah, the old gaming laptop debate rages on. I'm sure Jeremy has something to say about this too (maybe in his next column), but for me, I'm constantly flitting between a gaming laptop and a gaming desktop, but I'm currently in the latter camp.

I think it depends on your personal circumstances—as you mentioned, you're on the road a lot, and if I was traveling more, I would probably much rather spend $2,000 on a laptop than a desktop. But as I work from home most of the time, it's rare that I need to be lugging a laptop around with me, and on the rare occasions that I do, my work MacBook Pro or even Samsung Android tablet will do me just fine, especially the latter if I just want to play the odd game or catch up on some streaming shows.

That said, the latest generation of gaming laptops with Nvidia's 40-series GPUs have impressed me, ever since I saw them out at CES at the beginning of the year. The Samsung Galaxy Book3 Ultra (reviewed in the last issue) is a particular leap forward when it comes to design—it's as sleek and sexy as Apple's aforementioned pro laptop line, but also packing an RTX 4070 for some off-the-clock gaming. Okay, so it's not the best gaming experience compared to the RTX 4090 sitting in my desktop machine, but it will do fine for a quick game of *Forza Horizon 5* on the train.

So in summary, I think that gaming laptops are great, even if Jeremy's still on the fence about them. Have you used any good gaming laptops this year? If so, write in and let me know your recommendations.

OLED Dilemma

I read with interest your review of the Alienware AW3423DWF OLED gaming monitor, which you scored a 9. It's cheaper than its predecessor, the AW3423DW, but the older model supported NVIDIA G-SYNC Ultimate and a slightly higher 175Hz refresh rate. I've finally decided that now's the time to buy an OLED gaming monitor, but as an Nvidia RTX 4080 owner, I can't decide which one to buy! Help! **– J. Michaels**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: I'm so glad that you asked this, because I've used both these monitors myself. First of all, they're both easy to recommend, and you can't go wrong with either, so if it is a case of price or availability, just go for whichever is cheapest or easy to get.

That said, I personally prefer the AW3423DWF model that we reviewed in our August 2023 issue. It might have just been my model of the predecessor, but the cooling that's required as part of the G-SYNC module in the older monitor would drive me crazy, with a low but annoying fan whine as I worked during the day. It's fine if you wear headphones, but then again, if anyone else is in your space while it's on, I imagine that they would find it annoying, too. The newer model is still actively cooled, but the fans are silent 99 percent of the time.

The other reason to go for the newer model is because it's firmware upgradeable, meaning that issues can be addressed and features added over time. Again, because of the G-SYNC module, the predecessor cannot. As for when it comes to using G-SYNC on the newer model, I consulted some Reddit threads on how to calibrate the monitor correctly, and I cannot see any difference between the models at all—the Variable Refresh Rate from my RTX 4090 is silky smooth on both.

Enjoy joining the OLED fraternity! ⏻

THE BUILDS

THIS MONTH'S STREET PRICES...

BUDGET

HOW HIGH CAN you go before a budget rig stops being budget? That's a pertinent question with how GPU pricing is tumbling. On the AMD side, the Radeon RX 6800 vanilla GPU is now very cheap.

As we go to press, you can get one on Newegg for just $399. That's for a brand new card—not a refurb, not open box. What a deal. It's awfully tempting to slot it into the AMD budget rig and fiddle around with some other items to keep the price under $1,000. To be honest, that wouldn't be hard, given that with no changes, the total would come to $1,005.

But then that's a slippery slope. Why not find some serious savings throughout and try to drop in a 6800 XT for $495 and still keep the overall total below $1,000? With this in mind, we're not going to slap a $399 GPU into a budget rig, no matter how good a deal it seems. Happily, however, there's an even cheaper RX 6700 XT that we can stick in.

Okay, it's only $10 less than last month's PowerColor branded item. But it's still a bit cheaper, and it's likewise still pretty remarkable to find last-gen AMD cards that remain available, and just keep getting cheaper.

As for the Intel rig, the problem is that the gap between the 4060 and the 4060 Ti in price terms is pretty hefty. It's another $90 for the 4060 Ti. Intriguingly, you can now also have an RTX 3070 for $5 less than a 4060 Ti. That's tempting, because the 3070 is sometimes faster.

However, despite its wider 256-bit memory bus, the RTX 3070 remains an 8GB card. It also gives up support for some of the latest DLSS-powered features, including frame generation. Because of this, we're going to slam in a 4060 Ti this month and suck up that price increase—not that the 4060 Ti is a stellar deal, even at its newer, lower price. But it's quite a relief not to have to rely on that awful RTX 4060.

AMD INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$95
PSU	600W EVGA 600 GD 80+ Gold	$70
Mobo	MSI B550M PRO-VDH WIFI Micro ATX AM4 Motherboard	$105
CPU	AMD Ryzen 5 5600	$138
GPU	Sapphire Pulse AMD Radeon RX 6700 XT 12GB **NEW**	$309
RAM	16GB (2 x 8GB) PNY XLR8 Gaming EPIC-X RGB DDR4 @ 3600MHz	$45
SSD	1TB Lexar NM790 PCIe 4.0 M.2 SSD	$70
HDD	2TB Seagate Barracuda Compute 7200	$50
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$914**

INTEL INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$95
PSU	600W EVGA 600 GD 80+ Gold	$70
Mobo	ASRock B660M Pro RS Micro ATX GA1700	$95
CPU	Intel Core i5-12400F	$150
GPU	ZOTAC GAMING GeForce RTX 4060 Ti 8GB Twin Edge OC **NEW**	$369
RAM	16GB (2 x 8GB) PNY XLR8 Gaming EPIC-X RGB DDR4 @ 3600MHz	$45
SSD	1TB Lexar NM790 PCIe 4.0 M.2 SSD	$70
HDD	2TB Seagate Barracuda Compute 7200	$50
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$976**

MID-RANGE

WE GOT QUITE THE DEAL on the AMD Radeon 7800 XT last month. Sadly, prices have edged up since, returning to the card's $499 MSRP. That's quite a hike over the $430 we paid last issue.

Thus, we're going to drop the 7800 XT. Should we purloin the very cheap RX 6800 vanilla that we've just been toying with for the budget build? It's certainly one hell of a deal for just $399, and would pull our overall rig price down well under $1,500. On the other hand, the Radeon RX 6800 XT has been getting cheaper, too, and can be had for just $495 now.

That really is a very powerful GPU, and it's worth recalling that the new Radeon RX 7000 Series doesn't bring much to the table when it comes to new features. The biggest disappointment of all is that, dollar for dollar, the 7000 Series is barely any better for ray tracing than the 6000 Series, so it's a straight performance fight. The problem is that the 7800 XT is usually slightly faster, and even at $499 is only $5 more than the cheapest 6800 XT.

With that in mind, we've gone with the 6800 vanilla—a $100 saving, and a very appealing overall rig price. Moreover, the 6800 is close enough in performance to a 7800 XT that you'll rarely feel the performance gap, subjectively speaking.

As for the Intel rig, the good news is that Nvidia's RTX 4070 GPU just keeps getting cheaper. Last month, we had an MSI Ventus model for $550—down from the 5070's $599 MSRP. This month, that drops to $530 for a Zotac-branded 4070.

In truth, we'd prefer to see the 4070 drop under $500. At that point, there would be no doubting that it's a decent deal. We're not there quite yet, but we do seem to be heading in the right direction. If the rumors of a mid-life refresh for the RTX series are true, then the RTX 4070 and its Ada Lovelace-generation siblings will finally hit the prices that we think they should have been at launch.

AMD INGREDIENTS

PART		PRICE
Case	NZXT H7 Flow	$130
PSU	750W EVGA 750 BP 80+ Bronze	$90
Mobo	ASRock B650M PG RIPTIDE Wi-Fi Micro ATX AM5	$200
CPU	AMD Ryzen 5 7600X	$240
Cooler	MSI MAG CORELIQUID 240R V2	$90
GPU	XFX SPEEDSTER SWFT 319 Radeon RX 6800 **NEW**	$399
RAM	32GB (2x 16GB) Corsair Vengeance DDR5 CL40 @ 4800MHz	$98
SSD	1TB Lexar NM790 PCIe 4.0 M.2 SSD	$70
HDD	2TB Seagate Barracuda Compute 7200	$50
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$1,399**

INTEL INGREDIENTS

PART		PRICE
Case	NZXT H7 Flow	$130
PSU	750W EVGA 750 BP 80+ Bronze	$90
Mobo	Gigabyte Z790 UD AC ATX LGA 1700	$180
CPU	Intel Core i5-13600K	$317
Cooler	MSI MAG CORELIQUID 240R V2	$90
GPU	ZOTAC GAMING RTX 4070 Twin Edge DLSS 3 12GB **NEW**	$530
RAM	32GB (2x 16GB) Corsair Vengeance DDR5 CL40 @ 4800MHz	$98
SSD	1TB Lexar NM790 PCIe 4.0 M.2 SSD	$70
HDD	2TB Seagate Barracuda Compute 7200	$50
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$1,587**

TURBO

WITH NO NEW CPUS or GPUs from the big players of late, you can't expect to see dramatic changes in our Turbo builds. That said, the current GPU generation is very much maturing, and with that comes lower pricing than ever.

Last issue, the best price we could find for an AMD Radeon RX 7900 XT was $940. This time around, that price has dropped to $909. That's not far off $100 less than the 7900 XTX's original launch price. It's a no-brainer for our AMD build.

We're also happy to report that another $5 has been shaved off the price of the AMD Ryzen 9 7950X3D CPU. Okay, that's hardly the sort of savings that's going to have Dr Lazarus appealing to Grabthar's Hammer, but every little helps, even at the high end. And $655 for the best CPU AMD makes, complete with 16 cores and a big chunk of 3D V-cache, feels like a very decent deal.

Believe it or not, there are also now some slim savings to be had on the mighty Nvidia GeForce RTX 4090. The ZOTAC Gaming GeForce RTX 4090 Trinity 24GB is yours for $1,552 instead of the $1,599 MSRP.

In yet more good news, the Intel Core i9-13900KF is also very slightly cheaper this month. Okay, $9 is a fairly academic discount in the context of near-$3,500 PC. But again, it's indicative of the current market. We're in that lull between major component launches where the only really interesting developments involve pricing.

Oh, and if you want more price cuts, you can have another $35 off, this time courtesy of the Gigabyte Z790 UD AC motherboard, which is now down to just $145. That really is a great price for a quality Z790 motherboard that comes complete with integrated WiFi.

Moreover, what with the world economy still not having fully recovered post-pandemic, prices for all of these components are only likely to get better from a buyer's perspective. Watch this space over the next few months, as it's likely that these rigs are only going to get cheaper.

AMD INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$190
PSU	1000W NZXT C1000 80+ Gold	$180
Mobo	ASRock X670E PG Lightning ATX AM5	$230
CPU	AMD Ryzen 9 7950X3D	$655
Cooler	NZXT Kraken X73 RGB 360mm	$300
GPU	SAPPHIRE PULSE Radeon RX 7900 XTX 24GB **NEW**	$909
RAM	64GB (2x 32GB) Corsair Vengeance DDR5 CL40 @ 5200MHz	$165
SSD	2TB WD Black SN850X M.2 PCIe 4.0	$99
HDD	6TB WD Blue 5400 HDD	$100
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $2,860

INTEL INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$190
PSU	1000W Corsair RM1000x Shift 80+ Gold	$210
Mobo	Gigabyte Z790 UD AC ATX LGA 1700	$145
CPU	Intel Core i9-13900KF	$536
Cooler	NZXT Kraken X73 RGB 360mm	$300
GPU	Zotac GeForce RTX 4090 GAMING Trinity OC 24GB	$1,552
RAM	64GB (2x 32GB) Corsair Vengeance DDR5 CL40 @ 5200MHz	$165
SSD	2TB WD Black SN850X M.2 PCIe 4.0	$99
HDD	6TB WD Blue 5400 HDD	$100
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $3,329

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Holiday issue following the December issue, by Future US LLC, 130 West 42nd Street, 7th Floor, New York, NY 10036. USA. Website: www.futureus.com. Future US LLC also publishes *Mac|Life*, and *PC Gamer*. Entire contents copyright 2023, Future US LLC. All rights reserved. Reproduction in whole or in part is prohibited. Future US LLC. is not affiliated with the companies or products covered in *Maximum PC*. Reproduction on the Internet of the articles and pictures in this magazine is illegal without the prior written consent of *Maximum PC*. Products named in the pages of *Maximum PC* are trademarks of their respective companies.

PRODUCED IN THE UNITED STATES OF AMERICA.

UNEARTH THE FULL STORY BEHIND THE LEGEND OF ZELDA

Explore every game in the series and what makes each one so great. If you're a fan, you'll enjoy the trip through time. If you're a newcomer, we offer you this sage advice: "It's dangerous to go alone, take this!"

OLD!
retro GAMER PRESENTS
100% UNOFFICIAL
THE STORY OF ZELDA
DISCOVER THE COMPLETE HISTORY OF HYRU
BRAND NEW
TEARS OF THE KINGDOM REVEALED
EVERY GAME IN THE SERIES EXPLORED
TIMELINES EXPLAINED
BEHIND-THE-SCENES ACCESS
Digital Edition

ON SALE NOW

FUTURE

Shop easily online at:

magazinesdirect.com

bit.ly/3xbOGXq

PC GAMER
EDITOR'S CHOICE
MAXIMUM PC
KICK ASS!

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM

Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside

The Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm.Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.